arXiv:1610.06519v1 [math.OC] 20 Oct 2016

Stabilized Sparse Scaling Algorithms for Entropy Regularized Transport Problems

Bernhard Schmitzer

October 21, 2016

Abstract

Scaling algorithms for entropic transport-type problems have become a very popular numerical method, encompassing Wasserstein barycenters, multi-marginal problems, gradient flows and unbalanced transport. However, a standard implementation of the scaling algorithm has several numerical limitations: the scaling factors diverge and convergence becomes impractically slow as the entropy regularization approaches zero. Moreover, handling the dense kernel matrix becomes unfeasible for large problems. To address this, we propose several modifications: A log-domain stabilized formulation, the well-known ε-scaling heuristic, an adaptive truncation of the kernel and a coarse-to-fine scheme. This allows to solve larger problems with smaller regularization and negligible truncation error. A new convergence analysis of the Sinkhorn algorithm is developed, working towards a better understanding of ε-scaling. Numerical examples illustrate efficiency and versatility of the modified algorithm.

Contents

1 Introduction **2**
1.1 Motivation and Related Work . . . 2
1.2 Contribution and Outline . . . 4
1.3 Notation and Preliminaries . . . 5

2 Entropy Regularized Transport-Type Problems and Diagonal Scaling Algorithms **7**
2.1 Transport-Type Problems . . . 7
2.2 Entropy Regularization and Diagonal Scaling Algorithms . . . 9

3 Stabilized Sparse Multi-Scale Algorithm **12**
3.1 Log-Domain Stabilization . . . 12
3.2 Epsilon-Scaling . . . 14
3.3 Kernel Truncation . . . 15
3.4 Multi-Scale Scheme . . . 18

4 Analogy between Sinkhorn and Auction Algorithm **21**
4.1 Auction Algorithm . . . 21
4.2 Asynchronous Sinkhorn Algorithm and Iteration Bound . . . 23
4.3 Stability of Dual Solutions . . . 27
4.4 Application to Epsilon-Scaling . . . 31

5 Numerical Examples **33**
5.1 Setup . . . 33
5.2 Efficiency of Enhanced Algorithm . . . 34
5.3 Versatility . . . 36

6 Conclusion **42**

A Additional Proofs **43**
A.1 Proof of Lemma 4.20 . . . 43
A.2 Proof of Theorem 4.16 . . . 45

1 Introduction

1.1 Motivation and Related Work

Applications of Optimal Transport. Optimal transport (OT) is a classical optimization problem dating back to the seminal work of Monge and Kantorovich (see monographs [53, 3, 43] for introduction and historical context). The induced Wasserstein distances lift a metric from a 'base' space X to probability measures over X. This is a powerful analytical tool, for example to study PDEs as gradient flows in Wasserstein space [29, 4]. With the increase of computational resources, OT has also become a popular numerical tool in image processing, computer vision and machine learning (e.g. [42, 40, 54, 15, 50, 51, 46, 36, 23, 25]).

Many ideas have been presented to extend Wasserstein distances to general non-negative measures. We refer to [31, 17, 35, 18] and references therein for some context. A transport-type distance for general multi-channel signals is proposed in [52].

Computational Optimal Transport. To this day, the computational effort to solve OT problems remains the main bottleneck in many applications. In particular large problems, or even multi-marginal problems, remain challenging both in terms of runtime and memory demand.

For the linear assignment problem and discrete transport problems there are (combinatorial) algorithms based on the finite dimensional linear programming formulation by Kantorovich, such as the Hungarian method [33], the auction algorithm [10], the network simplex [2] and more [27]. Typically, they work for (almost) arbitrary cost functions, but do not scale well for large, dense problems. On the other hand, there are more geometric solvers, relying on the polar decomposition [14], that tend to be more efficient. There is the famous fluid dynamic formulation by Benamou and Brenier [6], explicit computation of the polar decomposition [28], semi-discrete solvers [37], and solvers of Monge-Ampère equation [9, 8] among many others. However, these only work on very specific cost functions, most notably the squared Euclidean distance. In a compromise between efficiency and flexibility, several discrete coarse-to-fine solvers have been proposed that adaptively select sparse sub-problems [45, 38, 44].

Entropy Regularization for Optimal Transport. Entropy regularization of optimal transport is related to the Schrödinger problem and the lazy gas experiment [34]. In the limit of vanishing regularization, primal and dual solutions of the entropic OT problem converge to particular solutions of the unregularized problem. A very general study of this limit is given in [34]. [16] provides a simpler and direct analysis for the 2-Wasserstein distance on $\mathbb{R}^d$, the discrete case is investigated in [20].

It has been observed that entropy regularization is also useful numerically [32] as it allows to use the Sinkhorn scaling algorithm [48], also called iterative projective fitting procedure (IPFP), to obtain an approximate solution to OT. More recently, practical applicability of Sinkhorn's algorithm to large problems has been popularized by [22] and today it is widely used, e.g. [49, 41, 52]. The algorithm starts from a kernel matrix (induced by the OT cost function) and iteratively rescales its rows and columns with diagonal scaling factors. Throughout the algorithm only the scaling factors need to be modified. For sufficiently large regularization this method has several advantages: The algorithm is very simple (and apt for parallelized implementation on GPUs), and its runtime is relatively low. And, as with discrete algorithms, its feasibility does not depend critically on the cost function. Moreover, the technique has been extended to more general transport-type problems, such as Wasserstein barycenters and multi-marginal problems [7], gradient flows [39] and unbalanced transport problems [19], resulting in a family of Sinkhorn-like diagonal scaling algorithms.

Convergence Speed of Sinkhorn Algorithm. In [26] the convergence rate of the Sinkhorn algorithm is studied for positive kernel matrices, yielding a global linear convergence rate of the marginals in terms of Hilbert's projective metric. However, applied to entropy regularized optimal transport, the contraction factor tends to one exponentially, as the regularization approaches zero, leading to a weak and rather impractical bound. In [30] the local convergence rate of the Sinkhorn algorithm near the solution is examined, based on a linearization of the iterations. This bound is tighter and more accurately describes the behaviour of the algorithm close to convergence. But these estimates do not apply when one starts far from the optimal solution, which is the usual case for small regularization parameters. In [32] a qualitative comparison is made between the Sinkhorn algorithm and the auction algorithm. In particular the role of the entropy regularization parameter is related to the slack parameter ε of the auction algorithm and it is pointed out that convergence of both algorithms becomes slower, as these parameters approach zero (but small parameters are required for good approximate solutions). For the auction algorithm this can provably be remedied by ε-scaling, where the ε parameter is gradually decreased during optimization. Analogously, it is suggested to gradually decrease entropy regularization during the Sinkhorn algorithm to accelerate optimization. Consequently, in the following we will also refer to the entropy regularization parameter as ε and to the gradual reduction scheme as ε-scaling. No quantitative analysis of the convergence speed or ε-scaling for the Sinkhorn algorithm is given in [32].

Using entropy regularization to approximately solve the linear assignment problem is also studied in [47].

Limitations of Entropic Transport. Despite its considerable merits, there are some fundamental constraints to the naive entropy regularization approach. The entropy term introduces blur in the optimal coupling, increasing with regularization. Sometimes this blur may be beneficial, such as in machine learning related classification tasks [22]. But in the majority of applications, blur is considered a nuisance (e.g. it quickly smears distinct features in gradient flows), and one would like to run the scaling algorithm with as little regularization as possible. However, a standard implementation has some major numerical limitations, becoming increasingly severe as the regularization approaches zero. The diagonal scaling factors diverge in the limit of vanishing regularization, leading to numerical overflow and instabilities. Moreover, the algorithm requires an increasing number of iterations to converge. In practice this can often be remedied by ε-scaling, but its efficiency is not yet well understood theoretically. Therefore,

numerically this limit is difficult to reach. In addition, naively storing the dense kernel matrix requires just as much memory as storing the full cost matrix in standard linear programming solvers and multiplications with the kernel matrix become increasingly slow. Thus, smart tricks to avoid storing of and multiplication by the dense kernel matrix have been conceived, such as efficient Gaussian convolutions or approximation by a pre-factored heat kernel [49]. However, these remedies only work for rather particular (although relevant) problems, and do not solve the issues of blur and diverging scaling factors.

1.2 Contribution and Outline

Contribution. The contributions presented in this article are twofold. Throughout Sect. 3 we propose an enhanced variant of the classical Sinkhorn algorithm to remedy the issues of entropy regularized optimal transport discussed in Sect. 1.1. These adaptions also apply to more general scaling algorithms for other transport-type problems, as presented in [19].

(i) We define a log-domain formulation of the scaling algorithm which avoids numerical overflow of the diagonal scaling factors but largely preserves the efficient matrix multiplication structure. This allows to robustly run the algorithm at small regularization parameters.

(ii) The already well-known and widely used ε-scaling heuristic (e.g. [32]) is used to reduce the number of required iterations.

(iii) We sparsify the kernel matrix by adaptive truncation, to reduce memory demand and accelerate iterations. We quantify the error induced by truncation and propose a truncation scheme which reliably yields small error bounds that are easy to evaluate.

(iv) Finally, a multi-scale representation of the transport problem is considered. This serves two purposes: First, it allows for a more efficient computation of the truncated kernel. Second, combining a coarse-to-fine approach with simultaneous ε-scaling drastically reduces the number of variables during early stages of ε-scaling, without losing much precision.

We will point out how each modification builds on the previous ones (Remark 5.1). Combining all four allows to solve a wide family of large transport-type problems with significantly less runtime, memory and regularization, as compared to the naive algorithm.

In Sect. 4 we develop an alternative convergence analysis of the Sinkhorn algorithm, different from the Hilbert metric approach given in [26]. Our eventual goal is a better theoretical understanding of the ε-scaling heuristic for the Sinkhorn algorithm. For this we revisit the analogy between the Sinkhorn and the auction algorithm, discussed in [32], and make a more quantitative comparison. We introduce a slightly modified asymmetric variant of the Sinkhorn iterations, that more closely mimics the behaviour of the auction algorithm and adapt the classical complexity analysis of the auction algorithm to this variant. The obtained asymptotic bound is in good agreement with numerical experiments. We then prove stability of optimal dual solutions of entropy regularized OT under changes of the regularization parameter. This complements results of [20] and is a step towards proving the efficiency of ε-scaling for the Sinkhorn algorithm. We study a simple counter example to show that additional assumptions on the problem will be required for a full proof. To our knowledge, these are the first theoretical results towards ε-scaling for the Sinkhorn algorithm.

Remark 1.1 (Comparison with [19])**.** The framework of diagonal scaling algorithms for entropic transport-type problems, as summarized in Sect. 2, was introduced in [19]. [19] focusses on theoretical analysis of the algorithms in a continuous setting, establishing for example existence of

iterates and convergence. In this article we focus on efficient numerical implementation of discretized problems. For self-containedness, the log-domain stabilization was also described in [19], but it only becomes applicable to larger problems in combination with the other modifications of Sect. 3.

Remark 1.2 (Comparison with [26] and [47])**.** The projective Hilbert metric, used in [26] to measure convergence, is very different from typical numerical choices, such as L^1 or L^∞ marginal errors. Therefore, an iteration estimate based on [26] does not accurately reflect practical performance of the algorithm. The analysis of Sect. 4 seems to be more suitable for this aspect (see also Remark 4.11 and Sect. 5.2).

In [47] a *deformed Sinkhorn* iteration is introduced, which in our context corresponds to gradually decreasing ε during the iterations. Asymptotic convergence of the primal iterates to an unregularized optimizer is shown, when ε is decreased slowly enough [47, Thm. 3.6]. However, this estimate is based on the convergence analysis by [26] and therefore requires an exponential number of iterations. In addition, a log-domain stabilization of the Sinkhorn algorithm is proposed. Compared to our approach it does however involve more numerical overhead (see Sect. 3.1).

Remark 1.3 (Comparison with [44])**.** Intuitively, the kernel truncation and coarse-to-fine scheme of Algorithm 3.4 are adaptions of the ideas in [44] to the diagonal scaling algorithms, but the related analysis is rather different. On one hand, due to the 'true' sparsity of linear programming (LP) solutions, in [44] convergence to the exact optimal solution can be shown. Whereas the scaling algorithm only solves the approximate regularized problem (up to a virtually negligible truncation error), and in pathological cases it might not fully converge (Example 3.5). However, in our experiments non-convergence was no issue and the approximation quality can be considered sufficient (Fig. 4). We observe that the sparse LP solver typically requires less memory (see Fig. 3). On the other hand, the scaling algorithms are much easier to adapt to new cost functions, since no shielding neighbourhoods must be constructed, but mere (somewhat accurate) hierarchical lower bounds suffice. It is also much easier to generalize to other transport-type problems which is non-trivial for the linear programming solver. Moreover, the scaling algorithms seem to scale more favourably with problem size (Fig. 3). We believe that these advantages outweigh the difficulties in many practical applications.

Outline. The paper is organized as follows: The introduction is concluded with notational conventions and preliminaries (Sect. 1.3). Then, we start in Sect. 2 by recalling the framework for transport-type problems, as well as the scaling algorithms for their entropy regularized variants, as presented in [19]. The various proposed adaptions to the basic scaling algorithm are discussed in Sect. 3. The comparison between the Sinkhorn algorithm and the auction algorithm is presented throughout Sect. 4. A numerical study of the efficiency of the various modifications and comparison to the analysis of Sect. 4, as well as various examples to illustrate the versatility of the diagonal scaling algorithms are given in Sect. 5. We conclude in Sect. 6.

1.3 Notation and Preliminaries

Throughout this article, we will consider transport problems between two discrete finite spaces X and Y. For a discrete, finite space Z (typically X, Y or $X \times Y$) we identify functions and measures over Z with vectors in $\mathbb{R}^{|Z|}$, which we simply denote by $\mathbb{R}^Z$. For $v \in \mathbb{R}^Z$, $z \in Z$ we write $v(z)$ for the component of v corresponding to z (subscript notation is reserved for other

purposes). The standard Euclidean inner product is denoted by $\langle \cdot, \cdot \rangle$. The sets of vectors with positive and strictly positive entries are denoted by $\mathbb{R}_+^Z$ and $\mathbb{R}_{++}^Z$. The probability simplex over Z is denoted by $\mathcal{P}(Z)$. We write $\overline{\mathbb{R}} \stackrel{\text{def.}}{=} \mathbb{R} \cup \{-\infty, +\infty\}$ for the extended real line and $\overline{\mathbb{R}}^Z$ for the space of vectors with possibly infinite components.

For $a, b \in \mathbb{R}^Z$ the operators $\odot$ and $\oslash$ denote pointwise multiplication and division, e.g. $a \odot b \in \mathbb{R}^Z$, $(a \odot b)(z) \stackrel{\text{def.}}{=} a(z) \cdot b(z)$ for $z \in Z$. The functions exp and log are extended to $\mathbb{R}^Z$ by pointwise application to all components: $\exp(a)(z) \stackrel{\text{def.}}{=} \exp(a(z))$. We write $a \geq b$ if $a(z) \geq b(z)$ for all $z \in Z$, $a \geq 0$ if $a(z) \geq 0$ for all $z \in Z$ (and likewise for $\leq$, $>$ and $<$). For $a \in \mathbb{R}$, a_Z denotes the vector in $\mathbb{R}^Z$ with all entries being a. We write $\max a$ and $\min a$ for the maximal and minimal entry of a.

For $\mu \in \mathbb{R}^Z$ and a subset $A \subset Z$ we also use the notation $\mu(A) \stackrel{\text{def.}}{=} \sum_{z \in A} \mu(z)$, analogous to measures. We say $\mu \in \mathbb{R}^Z$ is absolutely continuous w.r.t. $\nu \in \mathbb{R}_+^Z$ and write $\mu \ll \nu$ when $[\nu(z) = 0] \Rightarrow [\mu(z) = 0]$. This is the discrete special case of absolute continuity for measures. The set $\operatorname{spt} \mu \stackrel{\text{def.}}{=} \{z \in Z : \mu(z) \neq 0\}$ is called support of μ.

The power set of Z is denoted by 2^Z.

For a subset $\mathcal{C} \subset \mathbb{R}^Z$ the indicator function of $\mathcal{C}$ over $\mathbb{R}^Z$ is given by

$$\iota_{\mathcal{C}} : \mathbb{R}^Z \to \overline{\mathbb{R}}, \qquad v \mapsto \begin{cases} 0 & \text{if } v \in \mathcal{C}, \\ +\infty & \text{else.} \end{cases}$$

In particular, for $v, w \in \mathbb{R}^Z$ one finds $\iota_{\{v\}}(w) = 0$ if $v = w$ and $+\infty$ otherwise. For $v \in \mathbb{R}^Z$ we introduce the short notation

$$\iota_{\leq v} : \mathbb{R}^Z \to \overline{\mathbb{R}}, \qquad w \mapsto \begin{cases} 0 & \text{if } w(z) \leq v(z) \text{ for all } z \in Z, \\ +\infty & \text{else.} \end{cases}$$

Moreover, we merely write ι_+ for $\iota_{\mathbb{R}_+^Z}$.

The projection matrices $\mathrm{P}_X \in \mathbb{R}^{X \times (X \times Y)}$ and $\mathrm{P}_Y \in \mathbb{R}^{Y \times (X \times Y)}$ are given by

$$\mathrm{P}_X(x, (x', y')) \stackrel{\text{def.}}{=} \begin{cases} 1 & \text{if } x = x', \\ 0 & \text{else.} \end{cases} \qquad \mathrm{P}_Y(y, (x', y')) \stackrel{\text{def.}}{=} \begin{cases} 1 & \text{if } y = y', \\ 0 & \text{else.} \end{cases} \tag{1.1a}$$

They act on some $\pi \in \mathbb{R}^{X \times Y}$ as follows:

$$(\mathrm{P}_X \, \pi)(x) = \sum_{y \in Y} \pi(x, y) = \pi(\{x\} \times Y), \qquad (\mathrm{P}_Y \, \pi)(y) = \sum_{x \in X} \pi(x, y) = \pi(X \times \{y\}). \tag{1.1b}$$

That is, they give the X and Y marginal in the sense of measures. Conversely, for some $v \in \mathbb{R}^X$, $w \in \mathbb{R}^Y$ we find

$$(\mathrm{P}_X^\top v)(x, y) = v(x), \qquad (\mathrm{P}_Y^\top w)(x, y) = w(y). \tag{1.1c}$$

Definition 1.4 (Kullback-Leibler Divergence)**.** For $\mu, \nu \in \mathbb{R}^Z$ the Kullback-Leibler divergence of μ w.r.t. ν is given by

$$\mathrm{KL}(\mu | \nu) \stackrel{\text{def.}}{=} \begin{cases} \sum_{\substack{z \in Z: \\ \mu(z) > 0}} \mu(z) \log\left(\frac{\mu(z)}{\nu(z)}\right) - \mu(Z) + \nu(Z) & \text{if } \mu, \nu \geq 0, \mu \ll \nu, \\ +\infty & \text{else.} \end{cases} \tag{1.2}$$

We write KL^* for the convex conjugate w.r.t. the first argument and find

$$\mathrm{KL}^*(\alpha|\nu) = \sum_{z\in Z} (\exp(\alpha(z)) - 1) \cdot \nu(z)\,. \tag{1.3}$$

Throughout the article the KL divergence plays a central role and is used on various different base spaces. Sometimes, when referring to the KL divergence on a space Z, we will add a subscript KL_Z for clarification.

Definition 1.5 (KL Proximal Step)**.** For a convex function $f : \mathbb{R}^Z \to \overline{\mathbb{R}}$ and a step size $\tau > 0$ the proximal step operator for the Kullback-Leibler divergence is given by

$$\mathrm{prox}_{1/\tau} f : \mathbb{R}^Z \to \mathbb{R}^Z, \qquad \mu \mapsto \operatorname*{argmin}_{\nu\in\mathbb{R}^Z} \left(\tfrac{1}{\tau}\,\mathrm{KL}(\nu|\mu) + f(\nu)\right)\,. \tag{1.4}$$

A unique minimizer exists, if there is some $\nu \in \mathbb{R}^Z$, $\nu \ll \mu$ such that $f(\nu) \neq \pm\infty$. Throughout this article we shall always assume that this is the case.

For Sect. 4 we require the following Lemma.

Lemma 1.6 (Softmax and Softmin)**.** *For a parameter $\varepsilon > 0$ and $a \in \mathbb{R}^Z$ let*

$$\mathrm{softmax}(a,\varepsilon) \overset{\text{def.}}{=} \varepsilon \log\left(\sum_{z\in Z} \exp(a(z)/\varepsilon)\right), \quad \mathrm{softmin}(a,\varepsilon) \overset{\text{def.}}{=} -\varepsilon \log\left(\sum_{z\in Z} \exp(-a(z)/\varepsilon)\right).$$

For ε, $\lambda > 0$ and a, $b \in \mathbb{R}^Z$ one has the relations

$$\max(a) \le \mathrm{softmax}(a,\varepsilon) \le \max(a) + \varepsilon \log|Z|, \tag{1.5a}$$
$$\min(a) - \varepsilon \log|Z| \le \mathrm{softmin}(a,\varepsilon) \le \min(a), \tag{1.5b}$$
$$\min(a-b) - \lambda \log|Z| \le \mathrm{softmax}(a,\varepsilon) - \mathrm{softmax}(b,\lambda) \le \max(a-b) + \varepsilon \log|Z|, \tag{1.5c}$$
$$\min(a-b) - \varepsilon \log|Z| \le \mathrm{softmin}(a,\varepsilon) - \mathrm{softmin}(b,\lambda) \le \max(a-b) + \lambda \log|Z|. \tag{1.5d}$$

Proof. The first line follows immediately from $0 \le \exp(a(z)/\varepsilon) \le \exp(\max a/\varepsilon)$. Line three then follows from $\min(a-b) \le \max(a) - \max(b) \le \max(a-b)$. The second and fourth line are implied by $\mathrm{softmin}(a,\varepsilon) = -\mathrm{softmax}(-a,\varepsilon)$. □

We assume that the reader has a basic knowledge of convex optimization, such as convex conjugation, Fenchel-Rockafellar duality and primal-dual gaps (see for example [13, 5]).

2 Entropy Regularized Transport-Type Problems and Diagonal Scaling Algorithms

2.1 Transport-Type Problems

In [19] a family of transport-type optimization problems with a common functional structure was introduced, encompassing standard optimal transport, unbalanced transport formulations and gradient flows. The general structure is given in the following definition.

Definition 2.1 (Generic Transport-Type Problem)**.** For two convex marginal functions $F_X : \mathbb{R}^X \to \overline{\mathbb{R}}$, $F_Y : \mathbb{R}^Y \to \overline{\mathbb{R}}$ and a cost function $c \in \overline{\mathbb{R}}^{X \times Y}$ the primal transport-type problem is given by:

$$\min_{\pi \in \mathbb{R}^{X \times Y}} E(\pi) \qquad \text{with} \qquad E(\pi) \stackrel{\text{def.}}{=} F_X(\mathrm{P}_X\, \pi) + F_Y(\mathrm{P}_Y\, \pi) + \langle c, \pi \rangle + \iota_+(\pi) \tag{2.1a}$$

The corresponding dual problem is given by:

$$\max_{(\alpha,\beta) \in (\mathbb{R}^X, \mathbb{R}^Y)} J(\alpha, \beta) \qquad \text{with} \quad J(\alpha, \beta) \stackrel{\text{def.}}{=} -F_X^*(-\alpha) - F_Y^*(-\beta) - \iota_{\leq c}(\mathrm{P}_X^\top\, \alpha + \mathrm{P}_Y^\top\, \beta) \tag{2.1b}$$

The indicator function $\iota_{\leq c}(\mathrm{P}_X^\top\, \alpha + \mathrm{P}_Y^\top\, \beta)$ denotes the classical optimal transport dual constraint $\alpha(x) + \beta(y) \leq c(x, y)$ for all $(x, y) \in X \times Y$ (see Section 1.3).

This structure can be extended to multiple couplings to describe barycenter and multi-marginal problems (see [7, 19] for details). The standard optimal transport problem is obtained as a special case.

Definition 2.2 (Standard Optimal Transport)**.** For marginals $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$ and a cost function $c \in \overline{\mathbb{R}}^{X \times Y}$ the standard optimal transport problem is obtained from Def. 2.1 by setting $F_X \stackrel{\text{def.}}{=} \iota_{\{\mu\}}$, $F_Y \stackrel{\text{def.}}{=} \iota_{\{\nu\}}$. The primal and dual functional are given by:

$$E(\pi) = \iota_{\{\mu\}}(\mathrm{P}_X\, \pi) + \iota_{\{\nu\}}(\mathrm{P}_Y\, \pi) + \langle c, \pi \rangle + \iota_+(\pi) \tag{2.2a}$$

$$J(\alpha, \beta) = \langle \alpha, \mu \rangle + \langle \beta, \nu \rangle - \iota_{\leq c}(\mathrm{P}_X^\top\, \alpha + \mathrm{P}_Y^\top\, \beta) \tag{2.2b}$$

The set

$$\Pi(\mu, \nu) \stackrel{\text{def.}}{=} \{\pi \in \mathbb{R}_+^{X \times Y} : \mathrm{P}_X\, \pi = \mu,\ \mathrm{P}_Y\, \pi = \nu\} \tag{2.3}$$

is called the *couplings* between μ and ν. If $E(\pi) < \infty$, then $\pi \in \Pi(\mu, \nu)$.

Let $\pi^\dagger$ and $(\alpha^\dagger, \beta^\dagger)$ be a pair of primal and dual optimizers. If the optimal value is finite, then the following relation holds (see e.g. [53, Thm. 5.10]):

$$[\pi^\dagger(x, y) > 0] \Rightarrow [\alpha^\dagger(x) + \beta^\dagger(y) = c(x, y)] \tag{2.4}$$

The general framework of Def. 2.1 also allows the formulation of unbalanced transport problems, where the hard marginal constraints of Def. 2.2 are replaced by soft constraints. This allows meaningful comparison between measures of different total mass. Such formulations were studied for example in [35] (see also [19] for more context). A particularly relevant case for the soft constraints is the Kullback-Leibler divergence.

Definition 2.3 (Unbalanced Optimal Transport with KL Fidelity)**.** For marginals $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$, a cost function $c \in \overline{\mathbb{R}}^{X \times Y}$ and a weight $\lambda > 0$ the unbalanced transport problem with KL fidelity is given by:

$$E(\pi) = \lambda \cdot \mathrm{KL}(\mathrm{P}_X\, \pi | \mu) + \lambda \cdot \mathrm{KL}(\mathrm{P}_Y\, \pi | \nu) + \langle c, \pi \rangle + \iota_+(\pi) \tag{2.5}$$

$$J(\alpha, \beta) = -\lambda \cdot \mathrm{KL}^*(-\alpha/\lambda) - \lambda \cdot \mathrm{KL}^*(-\beta/\lambda) - \iota_{\leq c}(\mathrm{P}_X^\top\, \alpha + \mathrm{P}_Y^\top\, \beta) \tag{2.6}$$

When $X = Y$ is a metric space with metric d, for $\lambda = 1$ and the cost function $c = d^2$, the square root of the optimal value of (2.5) yields the so called Gaussian Hellinger-Kantorovich (GHK) distance on $\mathbb{R}_+^X$, introduced in [35]. Similarly, for the cost function

$$c(x,y) \stackrel{\text{def.}}{=} \begin{cases} -\log\left([\cos(d(x,y))]^2\right) & \text{if } d(x,y) < \pi/2 \\ +\infty & \text{else.} \end{cases} \tag{2.7}$$

one obtains the Wasserstein-Fisher-Rao (WFR) distance (or Hellinger-Kantorovich distance), introduced independently and simultaneously in [31, 17, 35]. WFR is the length distance induced by GHK [35].

2.2 Entropy Regularization and Diagonal Scaling Algorithms

Now we apply entropy regularization to the above transport-type problems (see Sect. 1.1 for references) and replace the non-negativity constraint in (2.1a) by the Kullback-Leibler divergence. For this we need to select a suitable reference measure $\rho \in \mathbb{R}_+^{X\times Y}$. We then replace the term $\iota_+(\pi)$ in (2.1a) by $\varepsilon \cdot \mathrm{KL}(\pi|\rho)$, where $\varepsilon > 0$ is a regularization parameter. Then one typically 'pulls' the linear cost term into the KL divergence:

$$\langle c, \pi\rangle + \varepsilon\ \mathrm{KL}(\pi|\rho) = \varepsilon\ \mathrm{KL}(\pi|K) + \varepsilon \cdot \big(\rho(X\times Y) - K(X\times Y)\big)$$

where $K \in \mathbb{R}_+^{X\times Y}$ with

$$K(x,y) \stackrel{\text{def.}}{=} \exp(-c(x,y)/\varepsilon) \cdot \rho(x,y)\,. \tag{2.8}$$

with the convention $\exp(-\infty) = 0$. K is called the kernel associated with c and the regularization parameter ε. For convenience we formally introduce the function

$$\mathrm{get}K : \mathbb{R}_{++} \to \mathbb{R}^{X\times Y}, \qquad \varepsilon \mapsto \exp(-c/\varepsilon) \odot \rho\,. \tag{2.9}$$

We obtain the regularized equivalent to Def. 2.1.

Definition 2.4 (Regularized Generic Formulation)**.**

$$\min_{\pi\in\mathcal{M}_+(X\times Y)} E(\pi) \qquad \text{with} \qquad E(\pi) \stackrel{\text{def.}}{=} F_X(\mathrm{P}_X\pi) + F_Y(\mathrm{P}_Y\pi) + \varepsilon\ \mathrm{KL}(\pi|K) \tag{2.10a}$$

$$\max_{(\alpha,\beta)\in(\mathbb{R}^X,\mathbb{R}^Y)} J(\alpha,\beta) \qquad \text{with} \quad J(\alpha,\beta) \stackrel{\text{def.}}{=} -F_X^*(-\alpha) - F_Y^*(-\beta) - \varepsilon\ \mathrm{KL}^*\left([\mathrm{P}_X^\top\alpha + \mathrm{P}_Y^\top\beta]/\varepsilon \middle| K\right) \tag{2.10b}$$

Primal optimizers $\pi^\dagger$ have the form

$$\pi^\dagger = \mathrm{diag}(\exp(\alpha^\dagger/\varepsilon))\ K\ \mathrm{diag}(\exp(\beta^\dagger/\varepsilon)) \tag{2.11}$$

where $(\alpha^\dagger, \beta^\dagger)$ are dual optimizers. Conversely, for dual optimizers $(\alpha^\dagger, \beta^\dagger)$, $\pi^\dagger$ constructed as above is primal optimal [19].

Intuitively we see the relation between (2.1) and (2.10) as $\varepsilon \to 0$. For example, the term $\varepsilon\ \mathrm{KL}^*\left([\mathrm{P}_X^\top\alpha + \mathrm{P}_Y^\top\beta]/\varepsilon \middle| K\right)$ in (2.10b) can be interpreted as a smooth barrier function for the dual constraint $\mathrm{P}_X^\top\alpha + \mathrm{P}_Y^\top\beta \le c$ in (2.1b). We refer to Sect. 1.1 for references to rigorous convergence results.

Remark 2.5 (Role of ρ)**.** In a continuous setup, the introduction of ρ is clearly necessary, as one requires a reference measure to take the KL divergence against. This role is sometimes neglected in articles on the discrete Sinkhorn algorithm and simply set to $\rho(x, y) = 1$. However, when discretizing an underlying continuous problem with increasingly higher resolution, it is important to take into account the consistent discretization of ρ: as the number of discrete bins increases, the mass of the coupling π will be split into more and more bins. If one keeps using $\rho(x, y) = 1$, the mass of the reference measure increases, and – with all other parameters (such as ε) being fixed – the regularization artifacts will increase with resolution. In unbalanced transport problems this will lead to an incentive to increase the total mass of the coupling.

Under suitable assumptions problem (2.10b) can be solved by alternating optimization in α and β (see [19] for details). For fixed β, consider the KL*-term:

$$\mathrm{KL}^*_{X\times Y}\left([\mathrm{P}_X^\top\alpha + \mathrm{P}_Y^\top\beta]/\varepsilon \middle| K\right) = \mathrm{KL}^*_X\left(\alpha/\varepsilon | K\,\exp(\beta/\varepsilon)\right) + \sum_{(x,y)\in X\times Y} K(x,y)\,\left(\exp(\beta(y)/\varepsilon) - 1\right).$$

Note that the last term is constant w.r.t. α. Therefore, optimizing (2.10b) over α, for fixed β corresponds to maximizing

$$J_X(\alpha) = -F^*_X(-\alpha) - \varepsilon\,\mathrm{KL}^*_X\left(\alpha/\varepsilon | K\,\exp(\beta/\varepsilon)\right), \tag{2.12}$$

where $K\,\exp(\beta/\varepsilon)$ denotes standard matrix vector multiplication. The corresponding primal problem consists of minimizing

$$E_X(\sigma) = F_X(\sigma) + \varepsilon\,\mathrm{KL}_X(\sigma | K\,\exp(\beta/\varepsilon)). \tag{2.13}$$

This is a proximal step of F_X for the KL divergence with step size $1/\varepsilon$ (see Def. 1.5). So, by using the PD-optimality conditions between (2.12) and (2.13) (see e.g. [5, Thm. 19.1]), for a given β the primal optimizer $\sigma^\dagger$ of (2.13) and the dual optimizer $\alpha^\dagger$ of (2.12) are given by

$$\sigma^\dagger = \mathrm{prox}_{\varepsilon F_X}(K\,\exp(\beta/\varepsilon)), \qquad \alpha^\dagger = \varepsilon\,\log(\sigma^\dagger \oslash (K\,\exp(\beta/\varepsilon))), \tag{2.14}$$

Analogously, partial optimization w.r.t. β for fixed α is related to KL proximal steps of F_Y. Starting from some initial $\beta^{(0)}$, we can iterate alternating optimization to obtain a sequence $\beta^{(0)}, \alpha^{(1)}, \beta^{(1)}, \alpha^{(2)}, \dots$ as follows:

$$\alpha^{(\ell+1)} \overset{\text{def.}}{=} \varepsilon\,\log\left(\mathrm{prox}_{\varepsilon F_X}(K\,\exp(\beta^{(\ell)}/\varepsilon)) \oslash [K\,\exp(\beta^{(\ell)}/\varepsilon)]\right), \tag{2.15a}$$

$$\beta^{(\ell+1)} \overset{\text{def.}}{=} \varepsilon\,\log\left(\mathrm{prox}_{\varepsilon F_Y}(K^\top\exp(\alpha^{(\ell+1)}/\varepsilon)) \oslash [K^\top\exp(\alpha^{(\ell+1)}/\varepsilon)]\right). \tag{2.15b}$$

The algorithm becomes somewhat simpler when it is formulated in terms of the effective variables

$$u \overset{\text{def.}}{=} \exp(\alpha/\varepsilon), \qquad v \overset{\text{def.}}{=} \exp(\beta/\varepsilon). \tag{2.16}$$

For more convenient notation we introduce the proxdiv operator of a function F and step size $1/\varepsilon$:

$$\mathrm{proxdiv}_{\varepsilon F}: \sigma \mapsto \mathrm{prox}_{\varepsilon F}(\sigma) \oslash \sigma \tag{2.17}$$

The iterations then become:

$$u^{(\ell+1)} \overset{\text{def.}}{=} \mathrm{proxdiv}_{\varepsilon F_X}(K\,v^{(\ell)}), \qquad v^{(\ell+1)} \overset{\text{def.}}{=} \mathrm{proxdiv}_{\varepsilon F_Y}(K^\top u^{(\ell+1)}). \tag{2.18}$$

This reduces the number of evaluations of logarithm and exponential function. The primal-dual relation (2.11) then becomes $\pi^\dagger = \mathrm{diag}(u^\dagger)\,K\,\mathrm{diag}(v^\dagger)$, which is why u and v are often referred to as diagonal scaling factors.

Remark 2.6. Throughout this article, we will refer to the arguments of the dual functionals (2.1b) and (2.10b) as *dual variables* and denote them with (α, β). The effective, exponentiated variables, introduced in (2.16), will be denoted by (u, v) and referred to as *scaling factors*.

For future reference let us state the full scaling algorithm.

Algorithm 2.1 (Scaling Algorithm)**.**

function ScalingAlgorithm(ε,$v^{(0)}$)
  $K \leftarrow \text{get}K(\varepsilon)$ // compute kernel, see (2.9)
  $v \leftarrow v^{(0)}$
  repeat
    $u \leftarrow \text{proxdiv}_{\varepsilon} F_X(K\, v)$
    $v \leftarrow \text{proxdiv}_{\varepsilon} F_Y(K^\top u)$
  until stopping criterion
  return (u, v)
end function

The stopping criterion is typically a bound on the primal-dual gap between dual iterates $(\alpha, \beta) = \varepsilon \log(u, v)$ and primal iterate $\pi = \text{diag}(u)\, K\, \text{diag}(v)$, an error bound on the marginals of π (for standard optimal transport) or a pre-determined number of iterations.

With alternating iterations (2.15) or (2.18) a large family of functionals of form (2.10a) can be optimized, as long as the KL proximal steps of F_X and F_Y can be computed efficiently. A particularly relevant sub-family is, where F_X and F_Y are separable and are a sum of pointwise functions. Then the KL steps decompose into pointwise one-dimensional KL steps, see [19, Section 3.4] for details.

Since Section 4 focusses on the special case of entropy regularized optimal transport, let us explicitly state the corresponding functional and iterations.

Definition 2.7 (Entropic Optimal Transport)**.** For marginals $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$ and a cost function $c \in \overline{\mathbb{R}}^{X \times Y}$ the entropy regularized optimal transport problem is obtained from Def. 2.4 by setting $F_X \overset{\text{def.}}{=} \iota_{\{\mu\}}$, $F_Y \overset{\text{def.}}{=} \iota_{\{\nu\}}$ (see Definition 2.2 for the unregularized functional). We find:

$$E(\pi) = \iota_{\{\mu\}}(\mathrm{P}_X\, \pi) + \iota_{\{\nu\}}(\mathrm{P}_Y\, \pi) + \varepsilon\ \mathrm{KL}(\pi | K) \tag{2.19a}$$

$$J(\alpha, \beta) = \langle \alpha, \mu \rangle + \langle \beta, \nu \rangle - \varepsilon\ \mathrm{KL}^*\left([\mathrm{P}_X^\top \alpha + \mathrm{P}_Y^\top \beta]/\varepsilon \middle| K \right) \tag{2.19b}$$

The proximal steps of F_X and F_Y are trivial (if K has non-empty columns and rows) and we recover the famous Sinkhorn iterations:

$$\text{proxdiv}_{\varepsilon} F_X(\sigma) = \mu \oslash \sigma\,, \qquad \text{proxdiv}_{\varepsilon} F_Y(\sigma) = \nu \oslash \sigma\,, \tag{2.20a}$$

$$u^{(\ell+1)} = \mu \oslash (K\, v^{(\ell)})\,, \qquad v^{(\ell+1)} = \nu \oslash (K^\top u^{(\ell+1)})\,. \tag{2.20b}$$

Remark 2.8 (Entropic OT and Alternating Projections)**.** Minimizing the regularized primal functional (2.19a) amounts to computing the KL projection of the kernel K onto the constraint set $\Pi(\mu, \nu)$, see (2.3). In addition to the interpretation as alternating maximization of the dual, the Sinkhorn iterations can also be viewed as alternating primal KL projections onto the row and column constraints. After each u and v iteration, we find respectively

$$\mathrm{P}_X\, \text{diag}(u^{(\ell+1)})\, K\, \text{diag}(v^{(\ell)}) = \mu\,, \qquad \mathrm{P}_Y\, \text{diag}(u^{(\ell+1)})\, K\, \text{diag}(v^{(\ell+1)}) = \nu\,.$$

For more details we refer to [7].

3 Stabilized Sparse Multi-Scale Algorithm

Throughout this section we propose four adaptions to the Algorithm 2.1, that overcome the limitations of a naive implementation outlined in Section 1.1.

3.1 Log-Domain Stabilization

When running Algorithm 2.1 with small regularization parameter ε, entries in the kernel K, and the scaling factors u and v may become both very small and very large, leading to numerical difficulties. However, under suitable conditions (see Section 1.1, in particular [20]), the optimal dual variables (α, β) remain finite and have a stable limit as $\varepsilon \to 0$. Therefore, we are looking for a reformulation of the algorithm, such that it remains numerically applicable with small ε.

In [47] it was proposed to rewrite the Sinkhorn iterations directly in terms of the dual variables, instead of the scaling factors. For example, an update of α would be performed as follows:

$$\psi^{(\ell+1)}(x,y) = -c(x,y) + \beta^{(\ell)}(y) \tag{3.1a}$$

$$\tilde{\psi}^{(\ell+1)}(x,y) = \psi^{(\ell+1)}(x,y) - \max_{y\in Y} \psi^{(\ell+1)}(x,y) \tag{3.1b}$$

$$\alpha^{(\ell+1)}(x) = \varepsilon \log \mu(x) - \varepsilon \log \left(\sum_{y\in Y} \exp(\tilde{\psi}^{(\ell+1)}(x,y)/\varepsilon) \cdot \rho(x,y) \right) - \max_{y\in Y} \psi^{(\ell+1)}(x,y) \tag{3.1c}$$

Subtracting the maximum from $\psi^{(\ell+1)}$ and adding it again later, avoids large arguments in the exponential function. While this resolves the issue of extreme scaling factors, it perturbs the simple matrix multiplication structure of the original Sinkhorn algorithm and requires many additional evaluations of the exponential function and the logarithm in each iteration.

As an alternative (as also presented in [19], see Remark 1.1) we employ a relative, redundant parametrization of the alternating dual optimization iterations. We write the scaling factors (u, v), (2.16), as

$$u = \tilde{u} \cdot \exp(\hat{\alpha}/\varepsilon)\,, \qquad\qquad v = \tilde{v} \cdot \exp(\hat{\beta}/\varepsilon)\,. \tag{3.2}$$

Our goal is to formulate iterations (2.18) directly in terms of $(\tilde{u}, \tilde{v})$, while keeping $(\hat{\alpha}, \hat{\beta})$ unchanged during most iterations. The role of $(\hat{\alpha}, \hat{\beta})$ is to occasionally 'absorb' the large values of (u, v) such that $(\tilde{u}, \tilde{v})$ remain bounded. This leads to two types of iterations: stabilized iterations, during which only $(\tilde{u}, \tilde{v})$ are changed, and absorption iterations, during which $(\tilde{u}, \tilde{v})$ are absorbed into $(\hat{\alpha}, \hat{\beta})$. In this way, we can combine the simplicity of the scaling algorithm in terms of the scaling factor formulation with the numerical stability of the iterations in the log-domain formulation.

Analogous to the function get$K$, (2.9), we define the *stabilized kernel* as

$$\mathrm{get}\mathcal{K} : \mathbb{R}^X \times \mathbb{R}^Y \times \mathbb{R}_{++} \to \mathbb{R}^{X\times Y}, \quad (\alpha, \beta, \varepsilon) \mapsto \mathrm{diag}(\exp(\alpha/\varepsilon))\, \mathrm{get}K(\varepsilon)\, \mathrm{diag}(\exp(\beta/\varepsilon)), \tag{3.3a}$$

$$[\mathrm{get}\mathcal{K}(\alpha, \beta, \varepsilon)](x,y) = \exp\left(-\tfrac{1}{\varepsilon}\left[c(x,y) - \alpha(x) - \beta(y)\right]\right) \cdot \rho(x,y)\,. \tag{3.3b}$$

Note that the second line, (3.3b), should be used to numerically evaluate the stabilized kernel, such that extreme values in (α, β) and c can cancel before exponentiation. With $K = \mathrm{get}K(\varepsilon)$ and $\mathcal{K} = \mathrm{get}\mathcal{K}(\hat{\alpha}, \hat{\beta}, \varepsilon)$ the following identities are readily verified:

$$K\, v = \exp(-\hat{\alpha}/\varepsilon) \odot \mathcal{K}\, \tilde{v}\,, \qquad\qquad K^\top u = \exp(-\hat{\beta}/\varepsilon) \odot \mathcal{K}^\top \tilde{u}\,. \tag{3.4}$$

Moreover, we introduce a stabilized version of the proxdiv operator:

$$\operatorname{proxdiv}_\varepsilon F : (\sigma, \gamma) \mapsto \operatorname{prox}_\varepsilon F(\exp(-\gamma/\varepsilon)\,\sigma) \oslash \sigma \tag{3.5}$$

Note that the regular version of the proxdiv operator, (2.17), is a special case of the stabilized variant with $\gamma = 0$. We observe that

$$\operatorname{proxdiv}_\varepsilon F(\mathcal{K}\,\tilde{v}, \hat{\alpha}) = \operatorname{proxdiv}_\varepsilon F(K\,v) \oslash \exp(\hat{\alpha}/\varepsilon), \tag{3.6a}$$

$$\operatorname{proxdiv}_\varepsilon F(\mathcal{K}^\top \tilde{u}, \hat{\beta}) = \operatorname{proxdiv}_\varepsilon F(K^\top u) \oslash \exp(\hat{\beta}/\varepsilon)\,. \tag{3.6b}$$

Now we formally state the stabilized variant of Algorithm 2.1.

Algorithm 3.1 (Stabilized Scaling Algorithm)**.**

1: **function** ScalingAlgorithmStabilized(ε,$\alpha^{(0)}$,$\beta^{(0)}$)
2: $(\hat{\alpha}, \hat{\beta}) \leftarrow (\alpha^{(0)}, \beta^{(0)})$
3: $(\tilde{u}, \tilde{v}) \leftarrow (1_X, 1_Y)$
4: $\mathcal{K} \leftarrow \text{get}\mathcal{K}(\hat{\alpha}, \hat{\beta}, \varepsilon)$
5: **repeat**
6: **while** $[\|\tilde{u}\|_\infty \le \tau] \wedge [\|\tilde{v}\|_\infty \le \tau]$ **do**
7: // stabilized iteration
8: $\tilde{u} \leftarrow \operatorname{proxdiv}_\varepsilon F_X(\mathcal{K}\,\tilde{v}, \hat{\alpha})$
9: $\tilde{v} \leftarrow \operatorname{proxdiv}_\varepsilon F_Y(\mathcal{K}^\top \tilde{u}, \hat{\beta})$
10: **end while**
11: // absorption iteration
12: $(\hat{\alpha}, \hat{\beta}) \leftarrow (\hat{\alpha}, \hat{\beta}) + \varepsilon \cdot \log(\tilde{u}, \tilde{v})$
13: $(\tilde{u}, \tilde{v}) \leftarrow (1_X, 1_Y)$
14: $\mathcal{K} \leftarrow \text{get}\mathcal{K}(\hat{\alpha}, \hat{\beta}, \varepsilon)$
15: **until** stopping criterion
16: $(\hat{\alpha}, \hat{\beta}) \leftarrow (\hat{\alpha}, \hat{\beta}) + \varepsilon \cdot \log(\tilde{u}, \tilde{v})$
17: **return** $(\hat{\alpha}, \hat{\beta})$
18: **end function**

Clearly, any subsequent application of stabilized iterations and absorption iterations in Algorithm 3.1 leads to an algorithm that is mathematically equivalent to Algorithm 2.1, in the sense that it produces the same iterates (keep in mind (3.2–3.6)). But numerically, with finite floating point precision, combining both types of iterations can make a significant difference. In practice one can run several stabilized iterations in a row, occasionally checking whether $(\tilde{u}, \tilde{v})$ become too large or too small (see line 6), and perform an absorption iteration if required. This inflicts less computational overhead than the direct log-domain formulation (3.1) and largely preserves the simple matrix multiplication structure of the scaling algorithms.

In the definitions for the stabilized kernel, (3.3b), and proxdiv-operator, (3.5), there still appear exponentials of the form $\exp(\cdot/\varepsilon)$, which may explode as $\varepsilon \to 0$. It is not immediately clear, if the modified 'stabilized' iterations are in fact numerically stable. We now address this question.

Let us first consider the case of standard optimal transport. Recall the primal interpretation of the Sinkhorn iterations as iterative KL projections onto the row and column constraints (Remark 2.8). Since the stabilized kernel is precisely the primal iterate at that step, it follows that its entries are bounded from above, and therefore so are the exponents in (3.3b). Entries

for which the exponents are extremely small may numerically be truncated to zero. However, as we will learn from Proposition 3.3, these truncation errors are negligible.

For the fixed marginal constraint $F_X = \iota_{\{\mu\}}$ we find that the stabilized proxdiv-operator (3.5) does not depend on γ and is equal to the standard proxdiv-operator of F_X (see Def. 2.7):

$$\operatorname{proxdiv}_\varepsilon F_X(\sigma, \gamma) = \operatorname{proxdiv}_\varepsilon F_X(\sigma) = \mu \oslash \sigma\,.$$

This invariance stems from the 'reparametrization invariance' of optimal transport, as described in the following Lemma. It implies that switching to the stabilized kernel corresponds to modifying the cost function in a way that does not affect the set of primal optimizers and merely adds an offset to the dual optimizers, which corresponds to the relative parametrization (3.2).

Lemma 3.1 (Reparametrization of Optimal Transport)**.** *For given marginals $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$, a cost function $c \in \overline{\mathbb{R}}^{|X|\times|Y|}$, and two functions $(\hat\alpha, \hat\beta) \in (\mathbb{R}^X, \mathbb{R}^Y)$ let $\hat c = c - \mathrm{P}_X^\top \hat\alpha - \mathrm{P}_Y^\top \hat\beta$. For the marginals μ, ν and cost functions c and $\hat c$, consider the corresponding standard optimal transport problems (Def. 2.2), and their regularized variants, Def. 2.7. Assume that the optimal values are finite.*

Then, $(\pi^\dagger, (\alpha^\dagger, \beta^\dagger))$ are primal and dual optimizers of the (un-)regularized problem for cost function c, if and only if $(\pi^\dagger, (\alpha^\dagger - \hat\alpha, \beta^\dagger - \hat\beta))$ are primal and dual optimizers for the (un-) regularized problem with modified cost $\hat c$.

Proof. If the primal-dual gap for $\pi^\dagger$ and $(\alpha^\dagger, \beta^\dagger)$ vanishes on the problem with cost function c, we find that it vanishes for $\pi^\dagger$ and $(\alpha^\dagger - \hat\alpha, \beta^\dagger - \hat\beta)$ on the problem with cost function $\hat c$, and vice versa. This hinges on the fact that the marginals of feasible $\pi^\dagger$ are fixed to μ and ν. □

Now we turn to more general transport-type problems, as laid out in Definition 2.4. For meaningful problems, F_X and F_Y will penalize unbounded growth of mass in the primal iterates π and Proposition 3.3 still holds. Consequently we can argue as above for the numerical robustness of the stabilized kernel. In the general case there is no invariance as given in Lemma 3.1 and the stabilized proxdiv-operator (3.5) depends on γ. In the examples studied in Section 5 and those given in [19] we find however, that evaluation of the exponential $\exp(-\gamma/\varepsilon)$ can be avoided.

It should be noted that one will face similar questions when trying to generalize the max-argument trick in (3.1) to more general scaling algorithms.

3.2 ε-Scaling

It is empirically and theoretically well-known (cf. Section 1.1) that convergence of Algorithm 2.1 becomes slow as $\varepsilon \to 0$. A popular heuristic remedy is the so-called ε-scaling, where one subsequently solves the regularized problem with gradually decreasing values for ε. Let $\mathcal{E} = (\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_n)$ be a list of decreasing positive parameters. We extend Algorithm 3.1 as follows:

Algorithm 3.2 (Scaling Algorithm with ε-Scaling)**.**

function ScalingAlgorithmεScaling($\mathcal{E}$,$\alpha^{(0)}$,$\beta^{(0)}$)
  $(\alpha, \beta) \leftarrow (\alpha^{(0)}, \beta^{(0)})$
  for $\varepsilon \in \mathcal{E}$ **do** // iterate over list, form largest to smallest
    $(\alpha, \beta) \leftarrow$ ScalingAlgorithmStabilized(ε,α,β)
  end for
  return (α, β)
end function

Note that the dual variable β is kept constant while changing ε, not the scaling factor v. This is because the optimal dual variables (α, β) usually have a stable limit as $\varepsilon \to 0$, while the scaling factors (u, v) typically diverge (see Sect. 1.1 and also Proposition 4.22). So far, virtually no theoretical results are available for ε-scaling of the Sinkhorn algorithm (e.g. Remark 1.2). We work towards a complexity analysis in Sect. 4.

3.3 Kernel Truncation

Storing the dense kernel K and computing dense matrix multiplications during the scaling iterations (2.18) requires a lot of memory and time on large problems. For several problems with particular structure, remedies have been proposed (Sect. 1.1). But these do not comprise non-standard cost functions, as the one used for the Wasserstein-Fisher-Rao distance, (2.7). Moreover they are not compatible with the log-stabilization (Section 3.1), thus a certain level of blur cannot be avoided. We are looking for a more flexible method to accelerate solving.

For many unregularized transport problems the optimal coupling $\pi^\dagger$ is concentrated on a sparse subset of $X \times Y$. In fact, this is the underlying mechanism for the efficiency of most solvers discussed in Section 1.1.

For the regularized problems the optimal coupling will usually be dense. This is due to the diverging derivative of the KL divergence at zero. However, as $\varepsilon \to 0$, the optimal coupling quickly converges to an unregularized solution (see Sect. 1.1, in particular [20, Thm. 5.8]). As $\varepsilon \to 0$, large parts of the coupling will approach zero exponentially fast.

So while we will not be able to exactly solve the full problem, by looking at suitable sparse sub-problems, we may still expect to obtain a reasonable approximation.

Let us formalize the concept of a sparse sub-problem.

Definition 3.2 (Sparse Sub-Problems)**.** Let F_X and F_Y be marginal functions and c be a cost function as in Definition 2.1 and let $\mathcal{N} \subset X \times Y$. We introduce:

$$\hat{c}(x,y) \stackrel{\text{def.}}{=} \begin{cases} c(x,y) & \text{if } (x,y) \in \mathcal{N}\,, \\ +\infty & \text{else.} \end{cases} \qquad \hat{K}(x,y) \stackrel{\text{def.}}{=} \begin{cases} K(x,y) & \text{if } (x,y) \in \mathcal{N}\,, \\ 0 & \text{else.} \end{cases} \tag{3.7}$$

We call problems (2.1a) and (2.1b) with c replaced by $\hat{c}$ the problems *restricted* to $\mathcal{N}$. This corresponds to adding the constraint $\operatorname{spt}\pi \subset \mathcal{N}$ to the primal problem, and only enforcing the constraint $\alpha(x) + \beta(y) \le c(x,y)$ on $(x,y) \in \mathcal{N}$ in the dual problem.

The entropy regularized variants of the restricted problems are obtained through replacing K by $\hat{K}$ in (2.10a) and (2.10b).

Clearly, when $\mathcal{N}$ is sparse, then so is $\hat{K}$ and the restricted regularized problem can be solved faster and with less memory. We now quantify the error inflicted by restriction.

Proposition 3.3 (Restricted Kernel and Duality Gap)**.** *Let $\varepsilon > 0$ and $\mathcal{N} \subset X \times Y$. Let E and J be unrestricted regularized primal and dual functionals with kernel K, as given in Definition 2.4, and let $\hat{E}$ and $\hat{J}$ be the functionals of the problems restricted to $\mathcal{N}$, with sparse kernel $\hat{K}$ (see Def. 3.2).*

Further, let (α, β) be a pair of dual variables, let $u = \exp(\alpha/\varepsilon)$, $v = \exp(\beta/\varepsilon)$ be the corresponding scaling factors and let $\pi = \operatorname{diag}(u)\, \hat{K}\, \operatorname{diag}(v)$ be the corresponding (restricted) primal coupling.

Then we find for the primal-dual gap between π and (α, β):

$$E(\pi) - J(\alpha,\beta) = \hat{E}(\pi) - \hat{J}(\alpha,\beta) + \sum_{(x,y)\in(X\times Y)\setminus\mathcal{N}} u(x)\, K(x,y)\, v(y)\,. \tag{3.8}$$

Proof. For the primal score we find:

$$E(\pi) = F_X(\mathrm{P}_X\,\pi) + F_Y(\mathrm{P}_Y\,\pi) + \varepsilon \sum_{(x,y)\in X\times Y} \left[\pi(x,y)\log\left(\tfrac{\pi(x,y)}{K(x,y)}\right) - \pi(x,y) + K(x,y)\right]$$

$$= \hat{E}(\pi) + \varepsilon \sum_{(x,y)\in (X\times Y)\setminus\mathcal{N}} \left[\underbrace{\pi(x,y)\log\left(\tfrac{\pi(x,y)}{K(x,y)}\right) - \pi(x,y)}_{=0} + K(x,y)\right]$$

Analogously, for the dual score we get:

$$J(\alpha,\beta) = -F_X^*(-\alpha) - F_Y^*(-\beta) - \varepsilon \sum_{(x,y)\in X\times Y} K(x,y)\cdot(\exp([\alpha(x)+\beta(y)]/\varepsilon) - 1)$$

$$= \hat{J}(\alpha,\beta) - \varepsilon \sum_{(x,y)\in (X\times Y)\setminus\mathcal{N}} K(x,y)\cdot\left(\underbrace{\exp([\alpha(x)+\beta(y)]/\varepsilon)}_{=u(x)\,v(y)} - 1\right)$$

Together we obtain:

$$E(\pi) - J(\alpha,\beta) = \hat{E}(\pi) - \hat{J}(\alpha,\beta) + \varepsilon \sum_{(x,y)\in (X\times Y)\setminus\mathcal{N}} K(x,y)\cdot(1 + u(x)\,v(y) - 1) \qquad \square$$

That is, the primal-dual gap for the original full functionals is equal to the gap for the truncated functionals plus the 'mass' that we have chopped off by truncating K to $\hat{K}$, when using the scaling factors u and v. If some $\mathcal{N}$ were known, on which most mass of the optimal $\pi^\dagger$ is concentrated, it would be sufficient to solve the problem restricted to $\mathcal{N}$, to get a good approximate solution. The remaining challenge is, how to identify $\mathcal{N}$ without knowing $\pi^\dagger$ before.

We propose an iterative re-estimation of $\mathcal{N}$, based on current dual iterates and to combine this with the log-stabilized iteration scheme (Section 3.1) and the computation of the stabilized kernel, (3.3b). For a threshold parameter $\theta > 0$ we define the following functions:

$$\mathrm{get}\mathcal{N}(\alpha,\beta,\varepsilon,\theta) \overset{\text{def.}}{=} \{(x,y)\in X\times Y:\ \exp(-\tfrac{1}{\varepsilon}[c(x,y)-\alpha(x)-\beta(y)]) \geq \theta\} \tag{3.9}$$

$$[\mathrm{get}\hat{K}(\alpha,\beta,\varepsilon,\theta)](x,y) \overset{\text{def.}}{=} \begin{cases} \exp(-\tfrac{1}{\varepsilon}[c(x,y)-\alpha(x)-\beta(y)])\,\rho(x,y) & \text{if } (x,y)\in \mathrm{get}\mathcal{N}(\alpha,\beta,\varepsilon,\theta)\,, \\ 0 & \text{else.} \end{cases} \tag{3.10}$$

The function $\mathrm{get}\hat{K}$ can be used instead of $\mathrm{get}\mathcal{K}$ in Algorithm 3.1. We refer to this as *absorption iteration with truncation*. For this combination one finds a simple bound for the primal-dual gap comparison of Proposition 3.3:

Proposition 3.4 (Simple Duality Gap Estimate for Absorption Iterations with Truncation)**.** *For a regularized problem as in Definition 2.4 with functionals E and J, let (u,v) be a pair of diagonal scaling factors and $(\alpha,\beta) = \varepsilon\cdot\log(u,v)$, let $(\hat{\alpha},\hat{\beta})$ a pair of dual variables and $(\tilde{u},\tilde{v})$ a pair of relative scaling factors such that*

$$u = \tilde{u}\cdot\exp(\hat{\alpha}/\varepsilon)\,, \qquad\qquad v = \tilde{v}\cdot\exp(\hat{\beta}/\varepsilon)\,.$$

Let further $\mathcal{N} = \mathrm{get}\mathcal{N}(\hat{\alpha},\hat{\beta},\varepsilon,\theta)$, $\mathcal{K} = \mathrm{get}\hat{K}(\hat{\alpha},\hat{\beta},\varepsilon,\theta)$, let $\hat{E}$ and $\hat{J}$ be the functionals restricted to $\mathcal{N}$ (see Definition 3.2) and let $\pi = \mathrm{diag}(\tilde{u})\,\mathcal{K}\,\mathrm{diag}(\tilde{v})$. Then

$$E(\pi) - J(\alpha,\beta) \leq \hat{E}(\pi) - \hat{J}(\alpha,\beta) + \|\tilde{u}\|_\infty\cdot\|\tilde{v}\|_\infty\cdot\theta\cdot\rho(X\times Y) \tag{3.11}$$

Proof. By virtue of Proposition 3.3

$$E(\pi) - J(\alpha,\beta) = \hat{E}(\pi) - \hat{J}(\alpha,\beta) + \sum_{(x,y)\in(X\times Y)\setminus\mathcal{N}} u(x)\,K(x,y)\,v(y)\,.$$

For $(x,y) \in (X \times Y) \setminus \mathcal{N}$ one has $\exp(-\frac{1}{\varepsilon}[c(x,y) - \hat{\alpha}(x) - \hat{\beta}(y)]) < \theta$ and therefore

$$\begin{aligned} u(x)\,K(x,y)\,v(y) &= \tilde{u}(x)\,\exp\Big(-\tfrac{1}{\varepsilon}[c(x,y) - \hat{\alpha}(x) - \hat{\beta}(y)]\Big)\cdot\rho(x,y)\cdot\tilde{v}(y) \\ &\le \tilde{u}(x)\,\tilde{v}(y)\,\theta\,\rho(x,y)\,. \end{aligned}$$

The result follows by bounding $\tilde{u}(x) \le \|\tilde{u}\|_\infty$, $\tilde{v}(y) \le \|\tilde{v}\|_\infty$ and summing over $(X \times Y) \setminus \mathcal{N}$. □

This implies that in Algorithm 3.1 with truncation the additional duality gap error due to the sparse kernel is bounded by

$$\|\tilde{u}^{(\ell)}\|_\infty \cdot \|\tilde{v}^{(\ell)}\|_\infty \cdot \theta \cdot \rho(X \times Y)\,.$$

In particular, before every stabilized iteration the error is bounded by $\tau^2 \cdot \theta \cdot \rho(X \times Y)$ and after every absorption iteration it is bounded by $\theta \cdot \rho(X \times Y)$. This bound is easy to evaluate and does not require to sum over $(X \times Y) \setminus \mathcal{N}$, as the exact expression in Proposition 3.3. We find that in practice this truncation error bound can be kept much smaller than the remaining primal-dual gap $\hat{E}(\pi) - \hat{J}(\alpha,\beta)$.

We point out that in general the stabilized iteration scheme *with truncation* might not converge. An illustration is given in Example 3.5. However, as we can infer from Proposition 3.4, if one regularly performs an absorption iteration before $\|\tilde{u}^{(\ell)}\|_\infty \cdot \|\tilde{v}^{(\ell)}\|_\infty$ becomes too large, the potential oscillations in the primal iterates and primal and dual functionals are numerically negligible.

Example 3.5 (Non-Convergence of Adaptive Sparse Scheme)**.** For some $\varepsilon > 0$, $C > 0$ let

$$\mu = \nu = \begin{pmatrix}\frac{1}{2} & \frac{1}{2}\end{pmatrix}^\top, \quad c = \begin{pmatrix}0 & C\\ C & 0\end{pmatrix}, \quad \rho = \begin{pmatrix}\frac{1}{2} & \frac{1}{2}\\ \frac{1}{2} & \frac{1}{2}\end{pmatrix}, \quad \alpha^{(0)} = \begin{pmatrix}0 & -C\end{pmatrix}^\top, \quad \beta^{(0)} = \begin{pmatrix}0 & C\end{pmatrix}^\top.$$

For $\theta \in (\exp(-2\,C/\varepsilon), 1]$ we find $\mathrm{get}\mathcal{N}(\alpha^{(0)},\beta^{(0)},\varepsilon,\theta) = \{(1,1),(1,2),(2,2)\}$. The corresponding $\mathrm{get}\hat{K}(\alpha^{(0)},\beta^{(0)},\varepsilon,\theta)$ does not have *total support* [48], since the entry $(1,2)$ does not lie on a positive diagonal. Consequently, the dual iterates diverge during the Sinkhorn algorithm. After ℓ iterations one obtains (up to a constant shift):

$$\alpha^{(\ell)} = \begin{pmatrix}0 & -C + \varepsilon\,\log(2\ell)\end{pmatrix}^\top, \qquad \beta^{(\ell)} = \begin{pmatrix}0 & C - \varepsilon\,\log(2\ell+1)\end{pmatrix}^\top.$$

So for $\ell \ge \frac{1}{2}\exp(2\,C/\varepsilon)$ the initial situation is approximately reversed and $\mathrm{get}\mathcal{N}(\alpha^{(\ell)},\beta^{(\ell)},\varepsilon,\theta) = \{(1,1),(2,1),(2,2)\}$. So for a suitable frequency of re-estimating $\mathcal{N}$, the algorithm can be made to oscillate between these two states. When $\mathcal{N}$ is re-estimated sufficiently often, oscillations in $\pi^{(\ell)}$ are small, and oscillations in $(\alpha^{(\ell)},\beta^{(\ell)})$ only reflect the degeneracy of the unregularized dual optimal set (cf. Example 4.21).

3.4 Multi-Scale Scheme

Finally, we propose to combine the stabilized sparse iterations with a hierarchical multi-scale scheme, analogous to the ideas in [37, 45, 38].

This serves two purposes: First, a hierarchical representation of the problem allows to determine the truncated sparse stabilized kernel get$\hat{K}$, (3.10), with a coarse-to-fine tree search, without explicitly testing all pairs $(x, y) \in X \times Y$. The second reason is to make the combination of ε-scaling (Algorithm 3.2) with the truncated stabilized scheme more efficient. For a fixed threshold θ, while ε is large, the support of the truncated kernel get$\hat{K}$ will contain many variables. At the same time, due to the blur induced by the regularization, the primal iterates will not provide a sharply resolved assignment. Solving the problems with large ε-value on a coarser grid reduces the number of required variables, without losing much spatial accuracy. As ε decreases, so does the number of variables in get$\hat{K}$ (since the exponential function decreases faster), and the resolution of X and Y can be increased. Therefore, it is reasonable to coordinate the reduction of ε with increasing the spatial resolution of the transport problem, until the desired regularization and resolution are attained.

We will now briefly recall the hierarchical representation of a transport problem from [45].

Definition 3.6 (Hierarchical Partition and Multi-Scale Measure Approximation [45])**.** For a discrete set X a *hierarchical partition* is an ordered tuple $(\mathcal{X}_0, \dots, \mathcal{X}_I)$ of partitions of X where $\mathcal{X}_0 = \{\{x\} \colon x \in X\}$ is the trivial partition of X into singletons and each subsequent level is generated by merging cells from the previous level, i.e. for $i \in \{1, \dots, I\}$ and any $\mathbf{x} \in \mathcal{X}_i$ there exists some $\hat{\mathcal{X}} \subset \mathcal{X}_{i-1}$ such that $\mathbf{x} = \bigcup_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}} \hat{\mathbf{x}}$. For simplicity we assume that the coarsest level is the trivial partition into one set: $\mathcal{X}_I = \{X\}$. We call $I > 0$ the *depth* of $\mathcal{X}$.

This implies a directed tree graph with vertex set $\bigcup_{i=0}^{I} \mathcal{X}_i$. For i, $j \in \{0, \dots, I\}$, $i < j$ we say $\mathbf{x} \in \mathcal{X}_i$ is a *descendant* of $\mathbf{x}' \in \mathcal{X}_j$ when $\mathbf{x} \subset \mathbf{x}'$. We call $\mathbf{x}$ a *child* of $\mathbf{x}$ for $i = j - 1$, and a *leaf* for $i = 0$.

For some $\mu \in \mathbb{R}^X$ its *multi-scale measure approximation* is the tuple $(\mu_0, \dots, \mu_I)$ of measures $\mu_i \in \mathbb{R}^{\mathcal{X}_i}$ defined by $\mu_i(\hat{\mathcal{X}}) = \mu(\bigcup_{\mathbf{x} \in \hat{\mathcal{X}}} \mathbf{x})$ for all subsets $\hat{\mathcal{X}} \subset \mathcal{X}_i$ and $i = 0, \dots I$.

For convenience we often identify X with the finest partition level $\mathcal{X}_0$, the set of singletons, and μ with μ_0.

Definition 3.7 (Hierarchical Dual Variables and Costs [45])**.** Let X and Y be discrete sets with hierarchical partitions $\mathcal{X} = (\mathcal{X}_0, \dots, \mathcal{X}_I)$, $\mathcal{Y} = (\mathcal{Y}_0, \dots, \mathcal{Y}_I)$ of depth I, let $\alpha \in \mathbb{R}^X$ and $\beta \in \mathbb{R}^Y$ be functions over X and Y, and let $c \in \mathbb{R}^{X \times Y}$ be a cost function.

Then we define the extension $\hat{\alpha} = (\hat{\alpha}_0, \dots, \hat{\alpha}_I)$ of α onto the full partition $\mathcal{X}$ by

$$\hat{\alpha}_i(\mathbf{x}) = \max_{x \in \mathbf{x}} \alpha(x) = \begin{cases} \alpha(x) & \text{if } i = 0 \text{ and } \mathbf{x} = \{x\} \text{ for some } x \in X, \\ \max_{\mathbf{x}' \in \text{children}(\mathbf{x})} \hat{\alpha}_{i-1}(\mathbf{x}') & \text{if } i > 0, \end{cases} \tag{3.12}$$

for $i \in \{0, \dots, I\}$ and $\mathbf{x} \in \mathcal{X}_i$ and analogous for $\hat{\beta}$ and β. Similarly, define an extension $\hat{c}$ of c by

$$\hat{c}_i(\mathbf{x}, \mathbf{y}) = \min_{(x,y) \in \mathbf{x} \times \mathbf{y}} c(x, y) \tag{3.13}$$

for $i \in \{0, \dots, I\}$, $\mathbf{x} \in \mathcal{X}_i$ and $\mathbf{y} \in \mathcal{Y}_i$.

For $i \in \{0, \dots, I\}$, $x \in \mathbf{x} \in \mathcal{X}_i$, $y \in \mathbf{y} \in \mathcal{Y}_i$ we find

$$\hat{c}_i(\mathbf{x}, \mathbf{y}) - \hat{\alpha}_i(\mathbf{x}) - \hat{\beta}_i(\mathbf{y}) \leq c(x, y) - \alpha(x) - \beta(y)\,. \tag{3.14}$$

Now we can implement a hierarchical tree-search for get$\mathcal{N}$ (and analogously get$\hat{K}$).

Algorithm 3.3 (Hierarchical Search for get$\mathcal{N}$).

function get$\mathcal{N}(\alpha,\beta,\varepsilon,\theta)$
  $(\hat{\alpha},\hat{\beta}) \leftarrow$ hierarchical extensions of (α,β) // see (3.12)
  $\mathcal{N} \leftarrow$ ScallCell$(\hat{\alpha},\hat{\beta},\varepsilon,\theta,I,\{X\},\{Y\})$ // call on coarsest partition level
end function

function ScanCell$(\hat{\alpha},\hat{\beta},\varepsilon,\theta,i,\mathbf{x},\mathbf{y})$
  $\mathcal{N}' \leftarrow \emptyset$ // temporary variable for result
  if $\hat{c}_i(\mathbf{x},\mathbf{y}) - \hat{\alpha}_i(\mathbf{x}) - \hat{\beta}_i(\mathbf{y}) \leq -\varepsilon \cdot \log\theta$ **then** // if cell cannot be ruled out at this level
    if $i > 0$ **then** // if not yet at finest level, check on all children
      for $(\mathbf{x}',\mathbf{y}') \in \text{children}(\mathbf{x}) \times \text{children}(\mathbf{y})$ **do**
        $\mathcal{N}' \leftarrow \mathcal{N}' \cup$ ScanCell$(\hat{\alpha},\hat{\beta},\varepsilon,\theta,i-1,\mathbf{x}',\mathbf{y}')$
      end for
    else // if at finest level, add variable
      $\mathcal{N}' \leftarrow \mathcal{N}' \cup (\mathbf{x} \times \mathbf{y})$ // recall $\mathbf{x} = \{x\}$, $\mathbf{y} = \{y\}$ for some $(x,y) \in X \times Y$ at $i = 0$
    end if
  end if
  return $\mathcal{N}'$
end function

From (3.14) follows directly that Algorithm 3.3 implements (3.9).

In many applications the discrete sets X and Y are point clouds in $\mathbb{R}^d$ and the hierarchical partitions are 2^d-trees over X and Y (see e.g. [44]). The cost function c is often originally defined on the whole product space $\mathbb{R}^d \times \mathbb{R}^d$ (such as the squared Euclidean distance). For the validity of Algorithm 3.3 it suffices if $\hat{c}_i(\mathbf{x},\mathbf{y}) \leq \min_{(x,y)\in\mathbf{x}\times\mathbf{y}} c(x,y)$. This allows to avoid computing (and storing) the full cost matrix $c \in \mathbb{R}^{X\times Y}$ and the explicit minimizations in (3.13). c and lower bounds on $\hat{c}_i$ can be computed on-demand directly using the tree-structure.

The second purpose of the multi-scale scheme is the combination with ε-scaling. As explained above, the purpose is to reduce the number of variables while ε is large. For an illustration see Fig. 1. For this, we divide the list $\mathcal{E}$ of regularization parameters ε into multiple lists $(\mathcal{E}_0,\dots,\mathcal{E}_I)$, with the largest values in $\mathcal{E}_I$ and the smallest (and final) values in $\mathcal{E}_0$, and sorted from largest to smallest within each $\mathcal{E}_i$. Then, for every i from I down to 0 we perform ε-scaling with list $\mathcal{E}_i$ at hierarchical level i, using the dual solution at each level as initialization at the next stage. The full algorithm, combining log-stabilization, ε-scaling, kernel truncation and the multi-scale scheme, is sketched next.

Algorithm 3.4 (Full Algorithm).

function ScalingAlgorithmFull$((\mathcal{E}_0,\dots,\mathcal{E}_I),\theta)$
  $i = I$
  $(\alpha,\beta) \leftarrow ((0),(0))$ // initialize dual variables
  while $i \geq 0$ **do**
    // solve problem at scale i with ε-scaling over $\mathcal{E}_i$
    for $\varepsilon \in \mathcal{E}_i$ **do** // iterate over list, from largest to smallest
      $(\alpha,\beta) \leftarrow$ ScalingAlgorithmStabilized$(i,\varepsilon,\theta,\alpha,\beta)$
    end for
    $i \leftarrow i - 1$
    if $i \geq 0$ **then** // refine dual variables
      $(\alpha,\beta) \leftarrow$ RefineDuals(i,α,β)

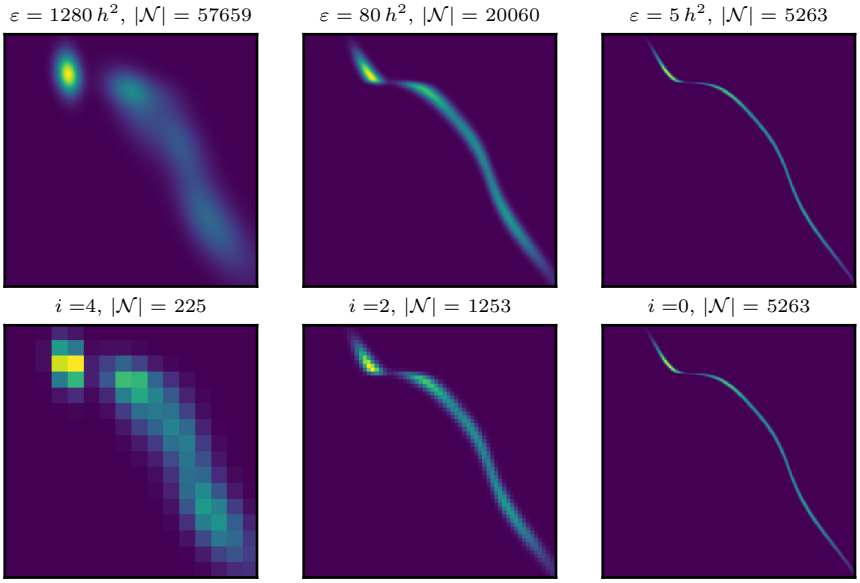
$\varepsilon = 1280\,h^2$, $|\mathcal{N}| = 57659$
$\varepsilon = 80\,h^2$, $|\mathcal{N}| = 20060$
$\varepsilon = 5\,h^2$, $|\mathcal{N}| = 5263$
i =4, $|\mathcal{N}| = 225$
i =2, $|\mathcal{N}| = 1253$
i =0, $|\mathcal{N}| = 5263$

Figure 1: ε-scaling, truncated kernels and multi-scale scheme. $X = Y$ is a uniform one-dimensional grid, representing $[0, 1]$, $|X| = 256$, $h = 256^{-1}$. μ and ν are smooth mixtures of Gaussians. *Top row* Density of optimal coupling $\pi^\dagger$ on X^2 for various ε. $|\mathcal{N}|$ is the number of variables in the truncated, stabilized kernel for fixed $\theta = 10^{-10}$. As ε decreases, so does $|\mathcal{N}|$, since $\pi^\dagger$ becomes more concentrated. *Bottom row* Optimal couplings for same ε as top row, but for different levels i of hierarchical partitions. i and ε were chosen to keep number of variables per $x \in X$ approximately constant. For high ε (and i) $|\mathcal{N}|$ is now dramatically lower. While $\pi^\dagger$ is 'pixelated' for high i, due to blur, it provides roughly the same spatial information as the top row. Images in third column are identical.

end if
end while
return (α, β)
end function

function RefineDuals(i,α,β)
  $(\alpha', \beta') \leftarrow (0_{\mathcal{X}_i}, 0_{\mathcal{Y}_i})$ // initialize refined duals with 0
  for $\mathbf{x} \in \mathcal{X}_i$ **do**
    $\alpha'(\mathbf{x}) \leftarrow \alpha(\text{parent}(\mathbf{x}))$
  end for
  for $\mathbf{y} \in \mathcal{Y}_i$ **do**
    $\beta'(\mathbf{y}) \leftarrow \beta(\text{parent}(\mathbf{y}))$
  end for
  return (α', β')
end function

Note: ScalingAlgorithmStabilized refers to calling Algorithm 3.1 for solving the problem at scale i, with get$\mathcal{K}$ replaced by get$\hat{K}$, (3.10), with threshold θ, implemented according to Algorithm 3.3. Accordingly, two arguments i and θ were added.

Remark 3.8 (Hierarchical Representation of F_X, F_Y)**.** To solve the problem at hierarchical scale

i, not only do we need a coarse version of c, as given in (3.13). In addition we need hierarchical versions of the marginal functions F_X, F_Y, see (2.10). An appropriate choice is often clear from the context of the problem. For example, for an optimal transport problem between μ and ν, see Def. 2.7, we set $F_{\mathcal{X}_i} = \iota_{\{\mu_i\}}$, where μ_i is taken from the multi-scale measure approximation of μ (see Def. 3.6). For the unbalanced transport problem with KL fidelity, Def. 2.3, we use $F_{\mathcal{X}_i} = \lambda \cdot \mathrm{KL}_{\mathcal{X}_i}(\cdot|\mu_i)$.

This completes the modifications of the diagonal scaling algorithm. Their usefulness will be demonstrated numerically in Sect. 5.

4 Analogy between Sinkhorn and Auction Algorithm

In this section we develop a new complexity analysis of the Sinkhorn algorithm and examine the efficiency of ε-scaling. To this end we will compare the Sinkhorn algorithm for the entropy regularized linear assignment problem with the auction algorithm. The similarity has already been pointed out in [32]. We will revisit this analogy and give a more quantitative comparison.

In this section we only consider the standard Sinkhorn algorithm (as opposed to general scaling algorithms), since the auction algorithm solves the linear assignment problem and assumptions on fixed marginals μ, ν are required for our analysis.

The auction algorithm is briefly recalled in Section 4.1. In Section 4.2 we introduce an asymmetric variant of the Sinkhorn algorithm, that is more similar to the original auction algorithm and provide an analogous worst-case estimate for the number of iterations until a given precision is achieved. A stability result for the dual optimal solutions under change of the regularization parameter ε is given in Section 4.3 and we discuss how it relates to ε-scaling in Sect. 4.4.

4.1 Auction Algorithm

For the sake of self-containedness, in this section we briefly recall the auction algorithm and its basic properties. Note that compared to the original presentation (e.g. [11]) we flipped the overall sign for compatibility with the notion of optimal transport.

In the following we consider a linear assignment problem, i.e. an optimal transport problem between two discrete sets X, Y with equal cardinality $|X| = |Y| = N$ where the marginals $\mu \in \mathbb{R}_+^X$, $\nu \in \mathbb{R}_+^Y$ are the counting measures. For simplicity we assume that the cost function $c \in \mathbb{R}^{X \times Y}$ is finite and non-negative.

The main loop of the auction algorithm is divided into two parts: During the bidding phase, elements of X that are unassigned determine their locally most attractive counterpart in Y (taking into account the current dual variables) and submit a bid for them. During the assignment phase, all elements of Y that received at least one bid, pick the most attractive one and change the current assignment accordingly. A formal description is given in the following.

Algorithm 4.1 (Auction Algorithm)**.**

function AuctionAlgorithm($\beta^{(0)}$)
  $\pi \leftarrow 0_{X \times Y}$ // initialize primal variable: 'empty' coupling
  $\beta \leftarrow \beta^{(0)}$
  while $\pi(X \times Y) < N$ **do**
    // bidding phase
    $B(y) \leftarrow \emptyset$ for all $y \in Y$ // initialize empty bid lists
    for $x \in \{x' \in X : \pi(\{x'\} \times Y) = 0\}$ **do** // iterate over unassigned x

8: $y \leftarrow \operatorname{argmin}_{y' \in Y}[c(x, y') - \beta(y')]$ // pick some element from argmin
9: $\alpha(x) \leftarrow c(x, y) - \beta(y)$ // set dual variable
10: $B(y) \leftarrow B(y) \cup \{x\}$ // submit bid to y, i.e. add x to bid list of y
11: **end for**
12: // assignment phase
13: **for** $y \in \{y' \in Y : B(y') \neq \emptyset\}$ **do** // iterate over all y that received bids
14: $\pi(\cdot, y) \leftarrow 0$ // set column of coupling to zero
15: $x \leftarrow \operatorname{argmin}_{x' \in B(y)}[c(x', y) - \alpha(x')]$ // find best bidder, pick one if multiple
16: $\beta(y) \leftarrow c(x, y) - \alpha(x) - \varepsilon$ // update dual variable
17: $\pi(x, y) \leftarrow 1$ // update coupling
18: **end for**
19: **end while**
20: **return** $(\pi, (\alpha, \beta))$
21: **end function**

Remark 4.1. In the above algorithm, line 9 is usually replaced by $\alpha(x) \leftarrow \min_{y' \in Y \setminus \{y\}}[c(x, y') - \beta(y')]$, which in practice may reduce the number of iterations. It does not affect the following worst-case analysis however, therefore we keep the simpler version.

We briefly summarize the main properties of the algorithm.

Proposition 4.2. *When $\varepsilon > 0$ and with dual initialization $\beta^{(0)} = 0_Y$, Algorithm 4.1 has the following properties:*

(i) α is increasing, β is decreasing.

(ii) After each assignment phase one finds:

$$\alpha(x) + \beta(y) \leq c(x, y)\,, \qquad [\pi(x, y) > 0] \Rightarrow [\alpha(x) + \beta(y) \geq c(x, y) - \varepsilon]\,. \tag{4.1}$$

The latter property is called ε-complimentary slackness. The primal iterate satisfies

$$\mathrm{P}_X \pi \leq \mu\,, \qquad \mathrm{P}_Y \pi \leq \nu\,. \tag{4.2}$$

(iii) The algorithm terminates after at most $N \cdot (C/\varepsilon + 1)$ iterations, where $C = \max c$.

Sketch of Proof. We note first, that $\beta(y)$ only changes upon accepting a bid, whereupon it decreases by at least ε, which implies that α is increasing. When $x \in X$ submits a bid to $y \in Y$ which gets accepted, we find directly that both properties (4.1) are satisfied and $\pi(x, y)$ is changed back to 0, before the next change in $\alpha(x)$ or $\beta(y)$ is made. When y accepts a bid from some other $x' \in X$, the dual constraint at (x, y) is still respected.

From the assignment phase we see that once $\pi(X \times \{y\}) = 1$ for some $y \in Y$, it remains 1. As long as $\pi(X \times \{y\}) = 0$, this y has not received a bid and therefore $\beta(y) = 0$. So while the algorithm has not yet terminated (and we execute the bidding phase), there must be at least one element $y^* \in Y$ with $\beta(y^*) = 0$.

If for some $y \in Y$ we have $\beta(y) < -C$, it will no longer receive bids, as it can no longer be minimal in line 8, since $c(x, y^*) - \beta(y^*) \leq C < c(x, y) - \beta(y)$. So the total number of accepted bids is bounded by $(C/\varepsilon + 1) \cdot N$. Since at least one bid is accepted by some $y \in Y$ per iteration, this bounds the total number of iterations.

For more details see for example [11]. □

From ε-complimentary slackness (4.1) we directly deduce the following result.

Corollary 4.3. *Upon convergence, the primal-dual gap of π and (α, β), cf. Def. 2.2, is bounded by*

$$\langle c, \pi \rangle - (\langle \mu, \alpha \rangle + \langle \nu, \beta \rangle) \leq N \cdot \varepsilon \,.$$

If c is integer and $\varepsilon < 1/N$, then the final primal coupling is optimal.

Remark 4.4 (Extension to Unbounded Cost Functions)**.** For cost functions that assign $c(x, y) = +\infty$ to some pairs, but still allow a feasible assignment, a more general estimate for number of decreases of β is given in [12, Lemma 5]. In the particular case of the auction algorithm this essentially leads to an additional factor N in the bound on the number of iterations.

Remark 4.5 (ε-Scaling for the Auction Algorithm)**.** During the auction algorithm it may happen that several elements in X compete for the same target $y \in Y$, leading to the minimal decrease of $\beta(y)$ by ε in each iteration. This phenomenon has been dubbed 'price haggling' [12] and can cause poor practical performance of the algorithm, close to the worst-case iteration bound. The impact of price haggling can be reduced by the ε-scaling technique, where the algorithm is successively run with a sequence of decreasing values for ε, each time using the final value of β as initialization of the next run (see also Algorithm 3.2). With this technique the factor C/ε in the iteration bound can essentially be reduced to a factor $\log(C/\varepsilon)$. An analysis of the ε-scaling technique for more general min-cost-flow problems can be found in [12].

4.2 Asynchronous Sinkhorn Algorithm and Iteration Bound

We now introduce a slightly modified variant of the standard Sinkhorn algorithm, derive an iteration bound and make a comparison with the auction algorithm. We emphasize that this modification is primarily made to facilitate theoretical study of the algorithm and do not advocate its merits in an actual implementation (cf. Remark 4.14).

For $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$ and a cost function $c \in \mathbb{R}_+^{X \times Y}$ we consider the entropic optimal transport problem (Def. 2.7). The reference measure ρ for regularization, see (2.8), is chosen to be the product measure $\rho(x, y) = \mu(x) \cdot \nu(y)$.

Let us state the modified Sinkhorn algorithm.

Algorithm 4.2 (Asynchronous Sinkhorn Algorithm)**.**

function AsynchronousSinkhorn(ε,$v^{(0)}$,q_{target})
  $K \leftarrow \text{get}K(\varepsilon)$ // compute kernel
  $v = v^{(0)}$
  repeat
    $u \leftarrow \mu \oslash (K\, v)$
    $\hat{v} \leftarrow \nu \oslash (K^\top u)$
    $v \leftarrow \min\{v, \hat{v}\}$ // element-wise minimum
    $\pi \leftarrow \text{diag}(u)\, K\, \text{diag}(v)$
    $q \leftarrow \pi(X \times Y)$
  until $q \geq q_{\text{target}}$
  return $(\pi, (u, v))$
end function

The only differences to the standard Sinkhorn algorithm (given by Algorithm 2.1 with proxdiv-operators (2.20)) lie in line 7 and in the choice of the specific stopping criterion (see Remark 4.11 for a discussion). In the standard algorithm one would set $v \leftarrow \hat{v}$. The modification implies that v is monotonously decreasing, which implies the following result for Algorithm 4.2 in the spirit of Proposition 4.2. This monotonicity is crucial for bounding the number of iterations.

Remark 4.6. Throughout this section, for clarity, we enumerate the iterates u, v, as well as the auxiliary variables $\hat{v}$, π and q in Algorithm 4.2, starting with the initial $v^{(0)}$ and proceeding with $(u^{(1)}, v^{(1)}, \hat{v}^{(1)}, \pi^{(1)}, q^{(1)}), \ldots$, similar to formulas (2.18). Moreover, we introduce the corresponding dual variable iterates $(\alpha^{(\ell)}, \beta^{(\ell)}, \hat{\beta}^{(\ell)}) = \varepsilon \cdot \log(u^{(\ell)}, v^{(\ell)}, \hat{v}^{(\ell)})$.

Proposition 4.7 (Monotonicity of Asynchronous Sinkhorn Algorithm)**.**

(i) *u and $\alpha = \varepsilon \log u$ are increasing, v and $\beta = \varepsilon \log v$ are decreasing, q is increasing.*

(ii) *$\mathrm{P}_X \pi \le \mu$ and $\mathrm{P}_Y \pi \le \nu$. We say π is* sub-feasible.

(iii) *There exists some $y^* \in Y$ such that $v(y^*) = v^{(0)}(y^*)$ for all iterations.*

Proof. By construction we have $v^{(\ell+1)} \le v^{(\ell)}$. Consequently $K\, v^{(\ell+1)} \le K\, v^{(\ell)}$ and thus $u^{(\ell+1)} \ge u^{(\ell)}$ and eventually $\hat{v}^{(\ell+1)} \le \hat{v}^{(\ell)}$.

After updating $u^{(\ell+1)}$ the row constraints are satisfied. That is $\mathrm{P}_X \operatorname{diag}(u^{(\ell+1)})\, K\, \operatorname{diag}(v^{(\ell)}) = \mu$. Since $v^{(\ell+1)}$ is only decreased (i.e. if the corresponding column constraint is violated from above), afterwards the iterate $\pi^{(\ell+1)}$ is sub-feasible.

Since $\hat{v}^{(\ell)}$ is decreasing, it follows that if $v^{(\ell)}(y) = \hat{v}^{(\ell)}(y)$ for some $y \in Y$, then $v^{(k)}(y) = \hat{v}^{(k)}(y)$ for all $k \ge \ell$. Let $Y^{(\ell)} = \{y \in Y : v^{(\ell)}(y) = \hat{v}^{(\ell)}(y)\}$. Then $Y^{(\ell)} \subset Y^{(\ell+1)}$. Conversely, if $y \notin Y^{(\ell)}$, then $v^{(\ell)}(y) < \hat{v}^{(\ell)}(y)$ and therefore $v^{(\ell)}(y) = v^{(0)}(y)$.

Let now $q^{(\ell)}(y) \overset{\text{def.}}{=} v^{(\ell)}(y)\, [K^\top u^{(\ell)}](y)$. If $y \in Y^{(\ell+1)}$, then $q^{(\ell+1)}(y) = \nu(y) \ge q^{(\ell)}(y)$ (as $q^{(\ell)}(y)$ can never exceed $\nu(y)$). If $y \notin Y^{(\ell+1)}$, then $v^{(\ell+1)}(y) = v^{(\ell)}(y) = v^{(0)}(y)$ and since $u^{(\ell)}$ is increasing, we find $q^{(\ell+1)}(y) \ge q^{(\ell)}(y)$. We obtain $q^{(\ell+1)} = \sum_{y \in Y} q^{(\ell+1)}(y) \ge q^{(\ell)}$.

When $Y^{(\ell)} \neq Y$, there exists some $y^* \in Y$ with $y^* \notin Y^{(k)}$, $v^{(k)}(y^*) = v^{(0)}(y^*)$ for $k \in \{1, \ldots, \ell\}$. If $Y^{(\ell)} = Y$, then $\hat{v}^{(\ell)} = v^{(\ell)} \le v^{(\ell-1)}$. By construction one has $(u^{(\ell)})^\top K\, v^{(\ell-1)} = \mu(X)$ and $(u^{(\ell)})^\top K\, \hat{v}^{(\ell)} = \nu(Y) = \mu(X)$. So if $Y^{(\ell)} = Y$, in fact $v^{(\ell)} = v^{(\ell-1)}$. Consequently, there exists some $y^* \in Y$ with $v^{(\ell)}(y^*) = v^{(0)}(y^*)$ for all iterations. □

Let us further investigate the increments of the dual variable iterates $\alpha^{(\ell)} = \varepsilon \log(u^{(\ell)})$.

Lemma 4.8 (Minimal Increment of $\alpha^{(\ell)}$)**.** *For $\ell \ge 1$ have $\langle \alpha^{(\ell+1)} - \alpha^{(\ell)}, \mu \rangle \ge \varepsilon\, (1 - q^{(\ell)})$.*

Proof. Recall that $\pi^{(\ell)} = \operatorname{diag}(u^{(\ell)})\, K\, \operatorname{diag}(v^{(\ell)})$, and introduce $\pi'^{(\ell)} = \operatorname{diag}(u^{(\ell+1)})\, K\, \operatorname{diag}(v^{(\ell)})$. Consider the following evaluations of the dual functional:

$$\begin{aligned} J(\alpha^{(\ell)}, \beta^{(\ell)}) &= \langle \alpha^{(\ell)}, \mu \rangle + \langle \beta^{(\ell)}, \nu \rangle - \varepsilon \cdot \pi^{(\ell)}(X \times Y) + \varepsilon \cdot K(X \times Y) \\ J(\alpha^{(\ell+1)}, \beta^{(\ell)}) &= \langle \alpha^{(\ell+1)}, \mu \rangle + \langle \beta^{(\ell)}, \nu \rangle - \varepsilon \cdot \pi'^{(\ell)}(X \times Y) + \varepsilon \cdot K(X \times Y) \end{aligned}$$

Note that $\pi^{(\ell)}(X \times Y) = q^{(\ell)}$, $\pi'^{(\ell)}(X \times Y) = 1$ and since going from $\alpha^{(\ell)}$ to $\alpha^{(\ell+1)}$ corresponds to a block-wise dual maximization have $J(\alpha^{(\ell+1)}, \beta^{(\ell)}) \ge J(\alpha^{(\ell)}, \beta^{(\ell)})$. The claim follows. □

With these tools we can bound the total number of iterations to reach a given precision.

Proposition 4.9 (Iteration Bound for the Asynchronous Sinkhorn Algorithm)**.** *Initializing with* $\beta^{(0)} = 0_Y \Leftrightarrow v^{(0)} = 1_Y$*, for a given* $q_{\text{target}} \in (0,1)$ *the number of iterations* n *necessary to achieve* $q^{(n)} \geq q_{\text{target}}$ *is bounded by*

$$n \leq 2 + \frac{C}{\varepsilon \cdot (1 - q_{\text{target}})}. \tag{4.3}$$

where $C = \max c$*. Moreover,* $\langle u^{(\ell)}, \mu \rangle \leq \exp(C/\varepsilon)$ *for all iterates* $\ell \geq 1$*.*

Proof. Let us look at the first 'bid' $\alpha^{(1)}$. With $c \geq 0$ we have

$$\begin{aligned} \alpha^{(1)}(x) &= \varepsilon \log\left(\frac{\mu(x)}{[K\, v^{(0)}](x)}\right) = \varepsilon \log\left(\frac{1}{\sum_{y\in Y} \nu(y) \cdot \exp(-c(x,y)/\varepsilon)}\right) \\ &\geq \varepsilon \log\left(\frac{1}{\sum_{y\in Y} \nu(y)}\right) = 0\,. \end{aligned} \tag{4.4}$$

By virtue of Proposition 4.7 $q^{(\ell)} \leq q^{(n)}$ for $\ell \leq n$. With Lemma 4.8 this implies

$$\langle \alpha^{(n)} - \alpha^{(1)}, \mu \rangle \geq \sum_{\ell=1}^{n-1} \varepsilon \cdot \underbrace{(1 - q^{(\ell)})}_{\geq 1 - q^{(n)}} \geq \varepsilon \cdot (n-1) \cdot (1 - q^{(n)})$$

and with (4.4)

$$\langle \alpha^{(n)}, \mu \rangle \geq \varepsilon \cdot (n-1) \cdot (1 - q^{(n)})\,. \tag{4.5}$$

From Proposition 4.7 we know that there is some $y^* \in Y$ with $v^{(\ell)}(y^*) = 1 \leq \hat{v}^{(\ell+1)}(y^*)$ for all iterates $\ell \geq 0$. So

$$\hat{v}^{(\ell+1)}(y^*) = \frac{\nu(y^*)}{[K^\top u^{(\ell+1)}](y^*)} = \frac{1}{\sum_{x\in X} \exp\left(-\frac{1}{\varepsilon}[c(x,y^*) - \alpha^{(\ell+1)}(x)]\right) \mu(x)} \geq 1\,,$$

from which we infer $\exp(-C/\varepsilon) \cdot \langle \exp(\alpha^{(\ell+1)}/\varepsilon), \mu \rangle \leq 1$, i.e. $\langle u^{(\ell+1)}, \mu \rangle \leq \exp(C/\varepsilon)$. With Jensen's inequality we eventually find $\langle \alpha^{(\ell+1)}, \mu \rangle \leq C$ for $\ell \geq 0$.

Combining this with (4.5) we obtain $n \leq 1 + \frac{C}{\varepsilon\,(1-q^{(n)})}$. So, as long as $q^{(n)} < q_{\text{target}}$ we have $n < 1 + \frac{C}{\varepsilon\,(1-q_{\text{target}})}$. By contraposition we know that there is some $n \leq 2 + \frac{C}{\varepsilon\,(1-q_{\text{target}})}$ such that $q^{(n)} \geq q_{\text{target}}$. □

And finally, we formally establish convergence of the iterates.

Corollary 4.10 (Convergence of Asymmetric Algorithm)**.** *The iterates* $(u^{(\ell)}, v^{(\ell)})$ *of the asymmetric Algorithm 4.2 converge to a solution of the scaling problem and* $q^{(\ell)} \to 1$*.*

Proof. With the upper bound $\langle u^{(\ell)}, \mu \rangle \leq \exp(C/\varepsilon)$ (Proposition 4.9) we obtain the pointwise lower bound $v^{(\ell)}(y) \geq \exp(-C/\varepsilon)$ for all $\ell \geq 0$. Since $v^{(\ell)}$ is pointwise decreasing, it converges to some limit $v^{(\infty)} \geq \exp(-C/\varepsilon) > 0$.

The map $f : v^{(\ell)} \mapsto v^{(\ell+1)}$ is continuous for $v^{(\ell)} > 0$. With $v^{(\ell)} \to v^{(\infty)}$ and $v^{(\ell+1)} = f(v^{(\ell)}) \to v^{(\infty)}$ have $f(v^{(\infty)}) = v^{(\infty)}$ which implies that $v^{(\infty)}$ (together with the corresponding $u^{(\infty)} = \mu \oslash (K\, v^{(\infty)})$) solves the scaling problem. This implies convergence of $q^{(\ell)}$ to 1. □

Remark 4.11 (On the Stopping Criterion)**.** The criterion $q \geq q_{\text{target}}$ is motivated by Lemma 4.8, to provide a minimal increment of α during iterations. $1-q$ measures the mass that is still missing and is equal to the L^1 error between the marginals of π and the desired marginals μ and ν. In pathological cases the dual variables (α, β) may still be far from optimizers, even though $q \geq q_{\text{target}}$ (see Example 4.12). In [26, Lemma 2] linear convergence of the marginals in the Hilbert projective metric is proven. This is a stricter measure of convergence, less prone to 'premature' termination. However, for small ε the contraction factor is roughly $1 - 4\exp(-C/\varepsilon)$, which is impractical. The scaling $\mathcal{O}(1/\varepsilon)$ predicted by Proposition 4.9 is consistent with numerical observations when one uses the L^1 or L^∞ marginal error as stopping criterion (Sect. 5.2). Therefore we consider the q-criterion to be a reasonable measure for convergence, as long as one keeps $1 - q_{\text{target}} \ll \delta$ (Example 4.12).

Example 4.12. We consider the 1×2 toy problem with the following parameters:

$$\mu = \begin{pmatrix} 1 \end{pmatrix}^\top, \qquad \nu = \begin{pmatrix} 1-\delta & \delta \end{pmatrix}^\top, \qquad c = \begin{pmatrix} 0 & C \end{pmatrix}, \qquad K = \begin{pmatrix} 1-\delta & \delta \cdot e^{-C/\varepsilon} \end{pmatrix}$$

for some $C > 0$, $\delta \in (0,1)$ and some regularization strength $\varepsilon > 0$. And we consider the scaling factors (one for X, two choices for Y):

$$u = \begin{pmatrix} 1 \end{pmatrix}^\top, \qquad v_1 = \begin{pmatrix} 1 & 1 \end{pmatrix}^\top, \qquad v_2 = \begin{pmatrix} 1 & e^{C/\varepsilon} \end{pmatrix}^\top.$$

Let $\pi_i = \operatorname{diag}(u)\, K\, \operatorname{diag}(v_i)$ and corresponding total masses q_i, $i = 1, 2$. We find:

$$\pi_1 = \begin{pmatrix} 1-\delta & \delta \cdot e^{-C/\varepsilon} \end{pmatrix}, \qquad q_1 = 1 - \delta\,(1 - e^{-C/\varepsilon}), \qquad \pi_2 = \begin{pmatrix} 1-\delta & \delta \end{pmatrix}, \qquad q_2 = 1.$$

π_2 and $(\alpha, \beta_2) = \varepsilon \log(u, v_2)$ are primal and dual solutions. π_1 is sub-feasible (see Proposition 4.7). For fixed $\varepsilon > 0$, as $\delta \to 0$, q_1 tends to 1 (but is strictly smaller), i.e. the pair (u, v_1) has almost converged in the q-measure sense, but the distance between $\beta_1 = \varepsilon \log v_1$ and the actual solution β_2 is C.

Remark 4.13 (Analogy to Auction Algorithm)**.** For now assume $|X| = |Y| = N$ and μ, ν are normalized counting measures. Then lines 5 and 6 in Algorithm 4.2, expressed in dual variables, become

$$\begin{aligned} \alpha(x) &\leftarrow \operatorname{softmin}(\{c(x,y) - \beta(y) | y \in Y\}, \varepsilon) + \varepsilon \log N\,, \\ \hat{\beta}(y) &\leftarrow \operatorname{softmin}(\{c(x,y) - \alpha(x) | x \in X\}, \varepsilon) + \varepsilon \log N\,. \end{aligned}$$

These are formally similar to the corresponding lines 9 and 16 in Algorithm 4.1. We can interpret line 5 in Algorithm 4.2 as x not just submitting a bid to the best candidate y, but to all candidates, weighted by the attractiveness (recall that in the Sinkhorn algorithm, a change in the dual variable directly implies a change in the primal iterate via (2.11)). Conversely, in line 7, y does not only accept the best bid, but bids from all candidates, again weighted by price. If there are too many bids (i.e. if the column constraint would be violated and $\hat{v}(y) < v(y)$), $\beta(y)$ decreases and thereby rejects superfluous offers.

Both algorithms are decidedly different from mere alternating optimizations of the unregularized dual transport problem (2.2b): In the auction algorithm there is an additional slack ε in the β-update, in the asynchronous Sinkhorn algorithm the min is replaced by softmin, as implied by the regularized dual (2.10b). These modifications are crucial for convergence of the algorithms.

In both algorithms the dual variables α and β (or the scaling factors u and v) are monotonous, as is the mass of the intermediate primal coupling π, which is not primal feasible during the

iterations (Propositions 4.2 and 4.7). In both algorithms there is a minimal step size for one of the dual variables ($-\varepsilon$ for β in the auction algorithm, see Lemma 4.8 for the asynchronous Sinkhorn algorithm). The minimal step size is proportional to ε, leading to an iteration bound that scales as $\mathcal{O}(\varepsilon^{-1})$. Thus, in both cases, a small value of ε is required to obtain accurate solutions (see Corollary 4.3 and Sect. 1.1), but the naive algorithms (without ε-scaling) require more iterations as ε decreases.

Two important differences are, that the iteration bound for the auction algorithm depends on N, while Proposition 4.9 does not. Moreover, the auction algorithm terminates after a finite number of steps, whereas the asynchronous Sinkhorn algorithm in general only converges asymptotically and the iteration bound diverges as $q_{\text{target}} \to 1$.

Remark 4.14 (Comparison with Symmetric Algorithm)**.** As mentioned above, Algorithm 4.2 can be turned into the standard Sinkhorn algorithm by replacing line 7 by $v \leftarrow \hat{v}$. The above analysis cannot be extended to this case, as the monotonicity properties of Proposition 4.7 are lost. But of course the symmetric algorithm still converges. While it cannot be proven that the symmetric iteration is always more efficient than the asymmetric one (there are very particular exceptions), in the overwhelming majority of practical cases, we observe that the symmetric algorithm is faster.

Despite these differences, the asymmetric variant and the above analysis provide some insight on why convergence of the Sinkhorn algorithm becomes slower as $\varepsilon \to 0$.

In the context of the comparison to the auction algorithm, one can interpret the standard (symmetric) Sinkhorn algorithm such that elements of y that do not receive sufficient bids (i.e. where $\hat{v} > v$), submit 'counter-bids' to X, to obtain mass more quickly. Such a symmetrization is in principle also possible for the auction algorithm. But then similarly the complexity analysis based on monotonous dual variables is no longer feasible, and the algorithm may even run indefinitely (see 'down iterations' in [12]).

4.3 Stability of Dual Solutions

The main result of this section is Theorem 4.16, which provides stability of dual solutions to entropy regularized optimal transport (Def. 2.7) under changes of the regularization parameter ε. Its implications for ε-scaling are discussed in Sect. 4.4.

We consider a similar setup as in Sect. 4.2: $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$, $c \in \mathbb{R}_+^{X \times Y}$. Again, the reference measure ρ for regularization, see (2.8), is chosen to be the product measure $\rho(x, y) = \mu(x) \cdot \nu(y)$. For Theorem 4.16 we introduce an additional assumption on μ and ν, its necessity is discussed in Example 4.21.

Assumption 4.15 (Atomic Mass)**.** For $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$ there is a positive integer $M < \infty$ such that $\mu = M^{-1} \cdot r$, $\nu = M^{-1} \cdot s$ for $r \in \mathbb{Z}_{++}^X$, $s \in \mathbb{Z}_{++}^Y$, where $\mathbb{Z}_{++}$ denotes the set of strictly positive integers.

Theorem 4.16 (Stability of Dual Solutions under ε-Scaling)**.** *Let* $\max\{|X|, |Y|\} \le N < \infty$ *and let* μ *and* ν *satisfy Assumption 4.15 for some positive integer* $M < \infty$*. For two regularization parameters* $\varepsilon_1 > \varepsilon_2 > 0$*, let* (α_1, β_1) *and* (α_2, β_2) *be maximizers of the corresponding dual regularized optimal transport problems (Def. 2.7) and let* $\Delta\alpha = \alpha_2 - \alpha_1$ *and* $\Delta\beta = \beta_2 - \beta_1$*. Then*

$$\max \Delta\alpha - \min \Delta\alpha \le \varepsilon_1 \cdot N \cdot (4 \log N + 24 \log M), \tag{4.6a}$$

$$\max \Delta\beta - \min \Delta\beta \le \varepsilon_1 \cdot N \cdot (4 \log N + 24 \log M). \tag{4.6b}$$

Note that the bounds in (4.6) do not depend on a bound on the cost function c.

Remark 4.17 (Proof Strategy)**.** The proof requires several auxiliary definitions and lemmas. The estimate consists of two contributions: One stems from following paths within connected components of what we call *assignment graph* (defined in the following Lemma), using the primal-dual relation (2.11). This reasoning is analogous to the proof strategy for ε-scaling in the auction algorithm (see [12]). However, between different connected components (2.11) is too weak to yield useful estimates. So a second contribution arises from a stability analysis of *effective diagonal problems* (in Lemmas 4.19 and 4.20).

Lemma 4.18 (Assignment Graph)**.** *For two feasible couplings π_1, $\pi_2 \in \Pi(\mu,\nu)$ and a threshold $M^{-1} \in \mathbb{R}$ the corresponding assignment graph is a bipartite directed graph with vertex sets (X, Y) and the set of directed edges*

$$\begin{aligned}\mathcal{E} = \{(x,y) \in X \times Y : \pi_2(x,y) \geq \mu(x)\cdot\nu(y)/M\} \\ \sqcup \{(y,x) \in Y \times X : \pi_1(x,y) \geq \mu(x)\cdot\nu(y)/M\}\end{aligned}$$

where $(a,b) \in \mathcal{E}$ indicates a directed edge from $a \to b$.

The assignment graph has the following properties:

(i) Every node has at least one incoming and one outgoing edge.

(ii) Let $X_0 \subset X$, $Y_0 \subset Y$ such that there are no outgoing edges from (X_0, Y_0) to the rest of the vertices, then $|\mu(X_0) - \nu(Y_0)| < 1/M$. This is also true when there are no incoming edges from the rest of the vertices. If μ and ν are atomic, with atom size $1/M$ (see Assumption 4.15), then $\mu(X_0) = \nu(Y_0)$.

(iii) Let μ and ν be atomic, with atom size $1/M$. Let $\{(X_i, Y_i)\}_{i=1}^R$ be the vertex sets of the strongly connected components of the assignment graph, for some $R \in \mathbb{N}$ (taking into account the orientation of the edges). Then the sets $\{X_i\}_{i=1}^R$ and $\{Y_i\}_{i=1}^R$ are partitions of X and Y, and $\mu(X_i) = \nu(Y_i)$ for $i = 1, \ldots, R$.

Proof. Assume, a node $x \in X$ had no outgoing edge. Then $\sum_{y\in Y} \pi_2(x,y) < \mu(x)/M \leq \mu(x)$. This contradicts $\pi_2 \in \Pi(\mu,\nu)$. Existence of incoming edges follows analogously.

Let $\hat{X}_0 = X \setminus X_0$, $\hat{Y}_0 = Y \setminus Y_0$. If (X_0, Y_0) has no outgoing edges, then

$$\begin{aligned}\sum_{(x,y)\in X_0\times\hat{Y}_0} \pi_2(x,y) < \sum_{(x,y)\in X_0\times\hat{Y}_0} \mu(x)\cdot\nu(y)/M \leq \frac{1}{M}\,, \\ \sum_{(x,y)\in \hat{X}_0\times Y_0} \pi_1(x,y) < \sum_{(x,y)\in \hat{X}_0\times Y_0} \mu(x)\cdot\nu(y)/M \leq \frac{1}{M}\,.\end{aligned}$$

Since π_1, $\pi_2 \in \Pi(\mu,\nu)$, the first line implies $\mu(X_0) < \nu(Y_0) + 1/M$ and the second line implies $\nu(Y_0) < \mu(X_0) + 1/M$, i.e. $|\mu(X_0) - \nu(Y_0)| < 1/M$. With Assumption 4.15 for atom size $1/M$, this implies $\mu(X_0) = \nu(Y_0)$. The statement about incoming edges follows from $\mu(\hat{X}_0) = 1 - \mu(X_0)$ and $\nu(\hat{Y}_0) = 1 - \nu(Y_0)$.

Every node in (X, Y) is part of at least one strongly connected component (containing at least the node itself). If two strongly connected components have a common element, they are identical. Hence, the strongly connected components form partitions of X and Y. For some $x \in X$ (or $y \in Y$), let $X_{\text{out}} \subset X$ and $Y_{\text{out}} \subset Y$ be the set of nodes that can be reached

from x, let $X_{\text{in}} \subset X$ and $Y_{\text{in}} \subset Y$ be the set of nodes from which one can reach x and let $(X_{\text{con}} = X_{\text{out}} \cap X_{\text{in}}, Y_{\text{con}} = Y_{\text{out}} \cap Y_{\text{in}})$ be the strongly connected component of x. Clearly $(X_{\text{out}}, Y_{\text{out}})$ has no outgoing edges. Hence, by item (ii) one has $\mu(X_{\text{out}}) = \nu(Y_{\text{out}})$. Moreover, $(X_{\text{out}} \setminus X_{\text{in}}, Y_{\text{out}} \setminus Y_{\text{in}})$ has no outgoing edges, hence $\mu(X_{\text{out}} \setminus X_{\text{in}}) = \nu(Y_{\text{out}} \setminus Y_{\text{in}})$, from which follows that $\mu(X_{\text{con}}) = \nu(Y_{\text{con}})$. □

Lemma 4.19 (Reduction to Effective Diagonal Problem)**.** *Let $\{X_i\}_{i=1}^R$ and $\{Y_i\}_{i=1}^R$ be partitions of X and Y, for some $R \in \mathbb{N}$, with $\mu(X_i) = \nu(Y_i)$ for $i = 1, \dots, R$. Let $\{y_i\}_{i=1}^R \subset Y$ such that $y_i \in Y_i$. Let $(\alpha^\dagger, \beta^\dagger)$ be optimizers for the dual entropy regularized optimal transport problem (Def. 2.7) for a regularization parameter $\varepsilon > 0$.*

Consider the following functional over $\mathbb{R}^R$:

$$\hat{J} : \mathbb{R}^R \to \mathbb{R}, \qquad \hat{\beta} \mapsto -\varepsilon \sum_{i,j=1}^{R} \exp\left(-\tfrac{1}{\varepsilon}\left[d(i,j) + \hat{\beta}(i) - \hat{\beta}(j)\right]\right)$$

where $d \in \mathbb{R}^{R \times R}$ with

$$d(i,j) = -\varepsilon \log\left(\sum_{\substack{x \in X_i \\ y \in Y_j}} \exp\left(-\tfrac{1}{\varepsilon}\left[c(x,y) - \alpha^\dagger(x) - \beta^\dagger(y_i) - \beta^\dagger(y) + \beta^\dagger(y_i)\right]\right) \cdot \mu(x) \cdot \nu(y)\right). \tag{4.7}$$

Then $\hat{\beta}^\dagger \in \mathbb{R}^R$, given by $\hat{\beta}^\dagger(i) = \beta^\dagger(y_i)$, is a maximizer of $\hat{J}$. Conversely, if $\hat{\beta}^{\dagger\dagger}$ is a maximizer of $\hat{J}$, then there is a constant $b \in \mathbb{R}$, such that $\hat{\beta}^{\dagger\dagger}(i) = \hat{\beta}^\dagger(i) + b$ for all $i \in 1, \dots, R$.

Proof. We define the functional $\hat{J} : \mathbb{R}^K \to \mathbb{R}$ as follows:

$$\hat{J} : \hat{\beta} \mapsto J\left(\begin{pmatrix}\tilde{\alpha}\\ \tilde{\beta}\end{pmatrix} + \begin{pmatrix}-B_X\\ B_Y\end{pmatrix}\hat{\beta}\right)$$

where J denotes the dual functional of entropy regularized optimal transport (2.19b), and

- $\tilde{\alpha} \in \mathbb{R}^X$ with $\tilde{\alpha}(x) = \alpha^\dagger(x) + \beta^\dagger(y_i)$ when $x \in X_i$;
- $\tilde{\beta} \in \mathbb{R}^Y$ with $\tilde{\beta}(y) = \beta^\dagger(y) - \beta^\dagger(y_i)$ when $y \in Y_i$;
- $B_X \in \mathbb{R}^{X \times R}$ with $B_X(x,i) = 1$ if $x \in X_i$ and 0 else;
- $B_Y \in \mathbb{R}^{Y \times R}$ with $B_Y(y,i) = 1$ if $y \in Y_i$ and 0 else.

Then one has

$$\begin{pmatrix}\alpha^\dagger\\ \beta^\dagger\end{pmatrix} = \begin{pmatrix}\tilde{\alpha}\\ \tilde{\beta}\end{pmatrix} + \begin{pmatrix}-B_X\\ B_Y\end{pmatrix}\hat{\beta}^\dagger\,.$$

Since maximizing $\hat{J}$ corresponds to maximizing J over an affine subspace, clearly $\hat{\beta}^\dagger$ is a maximizer of $\hat{J}$. Since $\hat{J}$ inherits the invariance of J under constant shifts, any $\hat{\beta}^{\dagger\dagger}$ of the form given above, is also a maximizer. Consequently, we may add the constraint $\hat{\beta}(1) = 0$, which does not change the optimal value. With this added constraint the functional becomes strictly convex, which implies a unique optimizer. Hence, any optimizer of the unconstrained functional can be written in the form of $\hat{\beta}^{\dagger\dagger}$.

Let us now give a more explicit expression of $\hat{J}(\hat{\beta})$. We find

$$\begin{aligned}\hat{J}(\hat{\beta}) = &\left\langle B_Y^\top \nu - B_X^\top \mu, \hat{\beta} \right\rangle \\ &- \varepsilon \sum_{i,j=1}^{R} \sum_{\substack{x \in X_i \\ y \in Y_j}} \exp\left(-\tfrac{1}{\varepsilon}\left[c(x,y) - \tilde{\alpha}(x) + \hat{\beta}(i) - \tilde{\beta}(y) - \hat{\beta}(j)\right]\right) \cdot \mu(x) \cdot \nu(y) \\ &+ \langle \mu, \tilde{\alpha} \rangle + \left\langle \nu, \tilde{\beta} \right\rangle + \varepsilon \cdot K(X \times Y)\,.\end{aligned}$$

Note that the third line is constant w.r.t. $\hat{\beta}$. Since $\mu(X_i) = \nu(Y_i)$ the linear term vanishes and we can write

$$\hat{J}(\hat{\beta}) = -\varepsilon \sum_{i,j=1}^{R} \exp\left(-\tfrac{1}{\varepsilon}\left[d(i,j) + \hat{\beta}(i) - \hat{\beta}(j)\right]\right) + \text{const}$$

with coefficients $d \in \mathbb{R}^{R \times R}$, as given above. The constant offset does not affect minimization. □

Lemma 4.20 (Effective Diagonal Problem and Stability)**.** *For a parameter $\varepsilon > 0$ and a real matrix $d \in \mathbb{R}^{R \times R}$ consider the following functional:*

$$\hat{J}_{\varepsilon,d}(\beta) = \sum_{i,j=1}^{R} \exp\left(\left[-d(i,j) - \beta(i) + \beta(j)\right]/\varepsilon\right) \tag{4.8}$$

Minimizers of $\hat{J}_{\varepsilon,d}$ exist.

Let $\varepsilon_1 > \varepsilon_2 > 0$ be two parameters and $d_1,\ d_2 \in \mathbb{R}^{R \times R}$ two real matrices. Let $\beta_1^\dagger$ and $\beta_2^\dagger$ be minimizers of $\hat{J}_{\varepsilon_1,d_1}$ and $\hat{J}_{\varepsilon_2,d_2}$, let $\Delta d = d_2 - d_1$, $\Delta\beta = \beta_2^\dagger - \beta_1^\dagger$. Let the matrix $w \in \mathbb{R}^{R \times R}$ be given by $w(i,j) = \max\{-\Delta d(i,j), \Delta d(j,i)\}$. Then

$$\max \Delta\beta - \min \Delta\beta \le \operatorname{maxdiam}(w) + 2\,\varepsilon_1\, R \log R\,,$$

where

$$\operatorname{maxdiam}(w) = \max\left\{ \sum_{i=1}^{k-1} w_{j_i,j_{i+1}} \;:\; k \in \{2,\dots,R\},\ j_i \in \{1,\dots,R\} \text{ for } i = 1,\dots,k,\ \text{all } j_i \text{ distinct.} \right\}.$$

That is, $\operatorname{maxdiam}(w)$ *is the length of the longest cycle-less path in $\{1,\dots,R\}$ with edge lengths w.*

The proofs of Theorem 4.16 and Lemma 4.20 can be found in Appendix A.

Example 4.21 (Necessity of Atomic Mass Assumption)**.** Assumption 4.15 is in fact necessary for Theorem 4.16. This can be illustrated by the following 2×2 example. Let

$$\mu = \begin{pmatrix} \frac{1}{2} + \delta & \frac{1}{2} - \delta \end{pmatrix}^\top, \qquad \nu = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \end{pmatrix}^\top, \qquad c = \begin{pmatrix} 0 & C \\ C & 0 \end{pmatrix}, \qquad K = \begin{pmatrix} 1 & e^{-C/\varepsilon} \\ e^{-C/\varepsilon} & 1 \end{pmatrix}$$

for a parameter $\delta \in [0, \frac{1}{2})$ and some regularization strength $\varepsilon > 0$. (For simplicity we have chosen the counting measure for ρ in K.) The dual optimizers of the unregularized problem are given by (to remove the ambiguity of constant shifts, we fix $\alpha(1) = 0$):

$$\alpha = \begin{pmatrix} 0 & \tau \end{pmatrix}^\top, \qquad \beta = \begin{pmatrix} 0 & -\tau \end{pmatrix}^\top,$$

where for $\delta \in (0, \frac{1}{2})$ one has $\tau = -C$, and for $\delta = 0$ the whole set $\tau \in [-C, C]$ is optimal. Denote by $(\alpha_\varepsilon, \beta_\varepsilon)$ the dual optimizers of the regularized problem with regularization parameter ε (where we fix again $\alpha_\varepsilon(1) = 0$ for uniqueness). It is possible, but tedious, to find explicit formulas for $(\alpha_\varepsilon, \beta_\varepsilon)$. Instead we give a qualitative discussion which we consider to provide more insight. According to [20, Prop. 3.2], as $\varepsilon \to 0$, $(\alpha_\varepsilon, \beta_\varepsilon)$ converges to the *centroid* of the set of dual optimizers. In our example, we find

$$\lim_{\varepsilon \to 0} \alpha_\varepsilon = \begin{pmatrix} 0 & \tau \end{pmatrix}^\top, \qquad \lim_{\varepsilon \to 0} \beta_\varepsilon = \begin{pmatrix} 0 & -\tau \end{pmatrix}^\top,$$

with $\tau = -C$ if $\delta > 0$ and $\tau = 0$ if $\delta = 0$. The limit changes abruptly, depending on the value of δ. In the regime $0 < \delta \ll e^{-C/\varepsilon} \ll 1$, the blur introduced by entropic smoothing is small compared to the scale of the cost function, but large compared to the small asymmetry introduced by $\delta > 0$. So $(\alpha_\varepsilon, \beta_\varepsilon)$ will be close to the $\tau = 0$ solution. But when one further decreases ε and comes into the regime $0 < e^{-C/\varepsilon} \ll \delta \ll 1$, one approaches the $\tau = -C$ solution. By making δ arbitrarily small (but non-zero), one can make the transition phase between the regimes for ε arbitrarily short, during which $\alpha_\varepsilon(2)$ changes by approximately C. Consequently, there can be no stability result independent of C for the dual solutions without quantizing δ.

4.4 Application To ε-Scaling

Assuming that we know the dual solution for some $\varepsilon_1 > 0$, then Theorem 4.16 allows bound the number of iterations of Algorithm 4.2 for some smaller $\varepsilon_2 \in (0, \varepsilon_1)$, independent of bounds on the cost function c.

Proposition 4.22 (Single ε-Scaling Step)**.** *Consider the set-up of Theorem 4.16. In particular, let $\varepsilon_1 > \varepsilon_2 > 0$ be two regularization parameters, let (α_1, β_1) and (α_2, β_2) be corresponding optimizers of (2.19b).*

If Algorithm 4.2 is initialized with $v^{(0)} = \exp(\beta_1/\varepsilon)$, for a given $q_{\text{target}} \in (0, 1)$ the number of iterations n necessary to achieve $q^{(n)} \geq q_{\text{target}}$ is bounded by

$$n \leq 2 + \frac{\varepsilon_1}{\varepsilon_2} \frac{N \cdot (4 \log N + 24 \log M) + \log M}{1 - q_{\text{target}}}. \tag{4.9}$$

Proof. For the optimal scaling factor u_1 of the ε_1-problem we find:

$$u_1(x) \overset{\text{def.}}{=} \exp(\alpha_1(x)/\varepsilon_1) = \left(\sum_{y \in Y} \exp\left(-\tfrac{1}{\varepsilon_1} [c(x,y) - \beta_1(y)] \right) \nu(y) \right)^{-1}$$

$$\Rightarrow u_1(x)^{-1} \, \nu(y)^{-1} \geq \exp\left(-\tfrac{1}{\varepsilon_1} [c(x,y) - \beta_1(y)] \right) \qquad \text{for all } (x,y) \in X \times Y$$

With this we can bound the first iterate of the ε_2-run of the algorithm by:

$$u^{(1)}(x) = \left(\sum_{y\in Y} \exp\left(-\tfrac{1}{\varepsilon_2}[c(x,y) - \beta_1(y)]\right)\nu(y)\right)^{-1}$$
$$\geq \left(\sum_{y\in Y} u_1(x)^{-\varepsilon_1/\varepsilon_2}\nu(y)^{-\varepsilon_1/\varepsilon_2}\,\nu(y)\right)^{-1} \geq u_1(x)^{\varepsilon_1/\varepsilon_2} M^{-\varepsilon_1/\varepsilon_2}$$

where we have used $\nu(y) \geq 1/M$, Assumption 4.15. Eventually we find

$$\alpha^{(1)}(x) \geq \alpha_1(x) - \varepsilon_1 \log M.$$

By monotonicity of the iterates we have $\beta_2 \leq \beta^{(\ell)} \leq \beta^{(0)} = \beta_1$ and $\beta^{(\ell)}(y') = \beta_1(y')$ for a suitable $y' \in Y$ (Proposition 4.2). Consequently $\max \Delta\beta = 0$. Then, from Theorem 4.16, we obtain $\beta_2(y) - \beta_1(y) \geq \min \Delta\beta \geq -\varepsilon_1 \cdot A$ where $A = N \cdot (4 \log N + 24 \log M)$. With this we can bound the u-iterates:

$$u^{(\ell)}(x) \leq u_2(x) \overset{\text{def.}}{=} \exp(\alpha_2(x)/\varepsilon_2) = \left(\sum_{y\in Y} \exp\left(-\tfrac{1}{\varepsilon_2}[c(x,y) - \beta_2(y)]\right)\nu(y)\right)^{-1}$$
$$\leq \left(\sum_{y\in Y} \exp\left(-\tfrac{1}{\varepsilon_2}[c(x,y) - \beta_1(y)]\right)\nu(y)\right)^{-1} \exp\left(\tfrac{\varepsilon_1}{\varepsilon_2} A\right)$$

With convexity of $s \mapsto s^{\varepsilon_1/\varepsilon_2}$ and Jensen's inequality we get

$$u^{(\ell)}(x) \leq \left(\sum_{y\in Y} \exp\left(-\tfrac{1}{\varepsilon_1}[c(x,y) - \beta_1(y)]\right)\nu(y)\right)^{-\varepsilon_1/\varepsilon_2} \exp\left(\tfrac{\varepsilon_1}{\varepsilon_2} A\right) = u_1(x)^{\varepsilon_1/\varepsilon_2} \exp\left(\tfrac{\varepsilon_1}{\varepsilon_2} A\right)$$

and finally

$$\alpha^{(\ell)}(x) \leq \alpha_1(x) + \varepsilon_1 A$$

Now using Lemma 4.8 and arguing as in Proposition 4.9, we find that there is some $n \leq 2 + \frac{\varepsilon_1}{\varepsilon_2}\frac{A+\log M}{1-q_{\text{target}}}$ such that $q^{(n)} \geq q_{\text{target}}$. □

Let now $C = \max c$ for a cost function $c \geq 0$, let $\hat{\varepsilon} > 0$ be the desired final regularization parameter, pick some $\lambda \in (0,1)$ and let $k \in \mathbb{N}$ such that $\hat{\varepsilon} \cdot \lambda^{-k} \geq C$. Let $\mathcal{E} = (\hat{\varepsilon} \cdot \lambda^{-k}, \hat{\varepsilon} \cdot \lambda^{-k+1}, \ldots, \hat{\varepsilon})$ be a list of decreasing regularization parameters.

Now we combine Algorithm 4.2 with ε-scaling, (cf. Algorithm 3.2). For $\varepsilon = \hat{\varepsilon} \cdot \lambda^{-k} \geq C$, according to Proposition 4.9 it will take at most $2 + \frac{1}{1-q_{\text{target}}}$ iterations. It is tempting to deduce from Proposition 4.22 that for each subsequent value of ε at most $2 + \frac{A}{\lambda(1-q_{\text{target}})}$ iterations are required, with $A = N \cdot (4 \log N + 24 \log M) + \log M$. For $N > 1$ the total number of iterations would then be bounded by $(2 + \frac{A}{\lambda(1-q_{\text{target}})}) \cdot (k+1)$. For fixed λ the step parameter k scales like $\log(C/\hat{\varepsilon})$. Consequently, the total number of iterations would be bounded by $\mathcal{O}(\log(C/\hat{\varepsilon}))$

w.r.t. the cost function and regularization, which would be analogous to ε-scaling for the auction algorithm (Remark 4.5).

There is an obvious gap in this reasoning: Theorem 4.16 assumes that β_1 is known exactly, while Algorithm 4.2 only provides an approximate result. From Example 4.12 we learn that in extreme cases this difference can be substantial and disrupt the efficiency of ε-scaling. Thus, additional assumptions on the problem are required to make the above argument rigorous.

However, as discussed in Remark 4.11, in practice we usually observe that approximate iterates are sufficient and we can therefore hope that ε-scaling does indeed serve its purpose.

5 Numerical Examples

Now we present a series of numerical experiments to confirm the usefulness of the modifications proposed in Sect. 3. We show that runtime and memory usage are reduced substantially. At the same time the adapted algorithm is still as versatile as the basic version of [19], Algorithm 2.1. But Algorithm 3.4 can solve larger problems at lower regularization, yielding very sharp results. We give examples for unbalanced transport, barycenters, multi-marginal problems and Wasserstein gradient flows. The code used for the numerical experiments is available from the author's website.[1]

5.1 Setup

We transport measures on $[0,1]^d$ for $d \in \{1,2,3\}$, represented by discrete equidistant Cartesian grids. The distance between neighbouring grid points is denoted by h. For the squared Euclidean distance cost function $c(x,y) = |x-y|^2$, x, $y \in \mathbb{R}^d$, K is a Gaussian kernel with approximate width $\sqrt{\varepsilon}$. Therefore, it is useful to measure ε in units of h^2. For $\varepsilon = h^2$ the blur induced by the entropy smoothing is on the length scale of one pixel. With the enhanced scaling algorithm we solve most problems in this section with $\varepsilon = 0.1 \cdot h^2$, leaving very little blur and giving a good approximation of the original unregularized problem (see Fig. 4).

Unless stated otherwise, we use the following settings: Test measures are mixtures of Gaussians, with randomized means and variances. The cost function is the squared Euclidean distance. ρ is the product measure $\mu \otimes \nu$ for optimal transport problems and the discretized Lebesgue measure on the product space for problems with variable marginals. For standard optimal transport the stopping criterion is the L^∞ error between prescribed marginals (μ, ν) and marginals of the primal iterate π (and likewise for Wasserstein barycenters). For all other models the primal-dual gap is used. We set $\theta = 10^{-20}$ for truncating the kernel and $\tau = 10^2$ as upper bound for $(\tilde{u}, \tilde{v})$ (cf. (3.10), Algorithm 3.1, line 6), implying a bound of $10^{-16} \cdot \rho(X \times Y)$ for the truncation error, which is many orders of magnitude below prescribed marginal accuracies or primal-dual gaps. The hierarchical partitions in the coarse-to-fine scheme are 2^d-trees, where each layer i is a coarser d-dimensional grid with grid constant h_i. For combination with ε-scaling (Algorithm 3.4) we choose the lists $\mathcal{E}_i$, $i > 0$, such that for the smallest ε_i in each $\mathcal{E}_i$ we have roughly $\varepsilon_i / h_i^2 \approx 1$. On the finest scale, we go down to the desired final value of ε. All reported run-times were obtained on a single core of an Intel Xeon E5-2697 processor at 2.7 GHz.

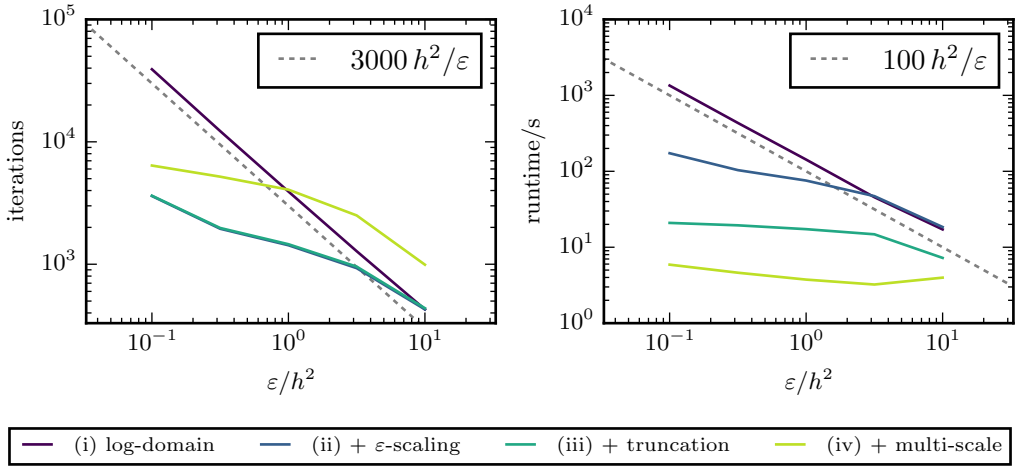

Figure 2: Efficiency of enhancements: average number of iterations and runtime for different ε and algorithms. $X = Y$ are 2-d 64×64 grids. (i) log-domain stabilized, Algorithm 3.1, (ii) with ε-scaling, Algorithm 3.2, (iii) with sparse stabilized kernel, (3.10), (iv) with multi-scale scheme, Algorithm 3.4. (ii) and (iii) need same number of iterations, but the sparse kernel requires less time. The naive implementation, Algorithm 2.1, requires same number of iterations as (i), but numerical overflow occurs at approximately $\varepsilon \leq 3\,h^2$.

5.2 Efficiency of Enhanced Algorithm

The numerical efficiency of the subsequent modifications presented in Sect. 3, applied to the standard Sinkhorn algorithm, is illustrated in Fig. 2. While the stabilized algorithm (i) is not yet faster than the naive implementation, it can robustly solve the problem for all given values of ε. The required number of iterations scales like $\mathcal{O}(1/\varepsilon)$, in good agreement with the complexity analysis of Sect. 4.2. With ε-scaling (ii) the number of iterations is decreased substantially. Replacing the dense kernel with the adaptive truncated sparse kernel (iii) does not change the number of required iterations, but saves time and memory. With the multi-scale scheme the required number of iterations is slightly increased, since the initial dual variables obtained at a coarser level are only approximate solutions. But by reducing the number of variables during the early ε-scaling stages, the runtime is further decreased (cf. Fig. 1). The combination of all modifications leads to an average total speed-up of more than two orders of magnitude on this problem type.

A runtime benchmark and study of the sparsity of the truncated kernel are given in Fig. 3. The runtime scales approximately linear with $|X|$ and for large problems the algorithm becomes faster than the adaptive sparse linear programming solver [44]. The final number of variables in the sparse kernel is comparable with the number of variables in [44], for higher values of ε, during scaling, more memory is required (cf. Fig. 4). This underlines again the importance of the coarse-to-fine scheme (Sect. 3.4). It should be noted, that Fig. 2 shows results for 64×64 images, the smallest image size in Fig. 3. For larger images the runtime difference between (i-iv) would be even larger, but due to time and memory constraints, only (iv) can be run practically.

The impact of different final values for ε is outlined in Fig. 4. As expected, the number of variables in the truncated kernel increases with ε. This leads to two competing trends in the

[1] http://wwwmath.uni-muenster.de/num/wirth/people/Schmitzer/

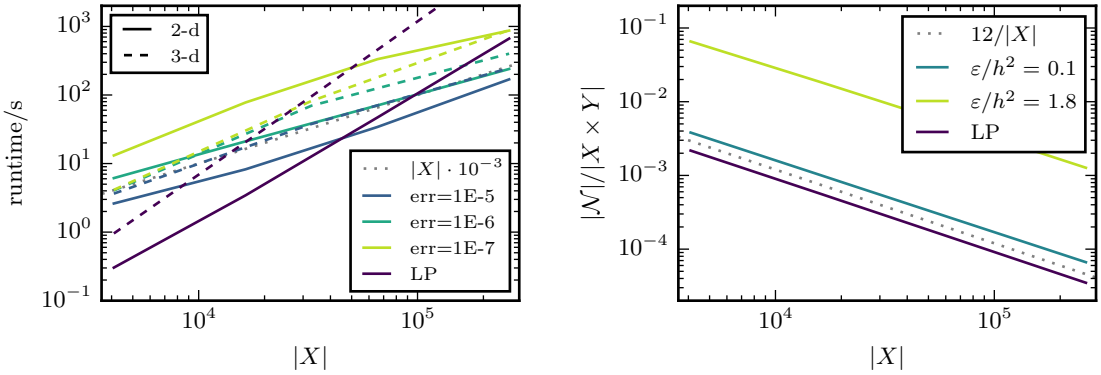

Figure 3: Average runtime and sparsity of Algorithm 3.4 for transporting test-images of different size (up to 512^2 pixels for 2-d, 64^3 for 3-d). Stopping criterion: L^∞-marginal error, for different accuracy limits, final $\varepsilon = 0.1 \cdot h^2$. Performance of the adaptive sparse linear programming solver [44] given for comparison (LP). As expected, runtime increases with required accuracy. The runtime of the scaling algorithm scales more favourably (approximately linear) with $|X|$ and is competitive for large instances. The number of variables scales as $\mathcal{O}(1/|X|)$, suggesting that the number of variables per $x \in X$ is roughly constant. For the final $\varepsilon = 0.1 \cdot h^2$, the sparsity of the truncated kernel is comparable to [44]. For $\varepsilon = 1.8 \cdot h^2$, the largest value in $\mathcal{E}_0$ (the list for the finest scale), more variables are required.

overall runtime: For large ε, the kernel truncation is less efficient, leading to an increase with ε. For small ε, the number of variables is very small, but more and more stages of ε-scaling are necessary, increasing the runtime as ε decreases further. Convergence of the regularized optimal dual variables to the unregularized optimal duals is exemplified in the right panel, justifying the use of the approximate entropy regularization technique for transport-type problems. While one may consider the dual sub-optimality at $\varepsilon \approx 30\, h^2$ sufficiently accurate, we point out that the corresponding primal coupling still contains considerable blur (cf. Fig. 1) and that due to less sparsity the runtime is actually higher than for $\varepsilon \approx h^2$.

As illustrated by Figs. 3 and 4, by choosing the threshold for the stopping criterion and the desired final ε, one can tune between required precision and available runtime.

Remark 5.1 (Interplay of Modifications)**.** The numerical findings presented in Figs. 2-4 underline how each of the modifications discussed in Sect. 3 builds on the previous ones and that all four of them are required for an efficient algorithm. The log-domain stabilization is an indispensable prerequisite for running the scaling algorithms with small regularization. However, for small ε, convergence tends to become extremely slow (cf. Fig. 2), therefore ε-scaling is needed to reduce the number of iterations. For small ε, kernel truncation significantly reduces the number of variables and accelerates the algorithm (cf. Figs. 2 and 4). However, for large ε (which must be passed during ε-scaling), far fewer variables are truncated and the algorithm cannot be run on large problems. This can be avoided by using the coarse-to-fine scheme, completing the algorithm. In principle it is possible, only to combine log-domain stabilization with kernel truncation, and to skip ε-scaling and the coarse-to-fine scheme. While this tends to solve the stability and memory issues, convergence is still impractically slow.

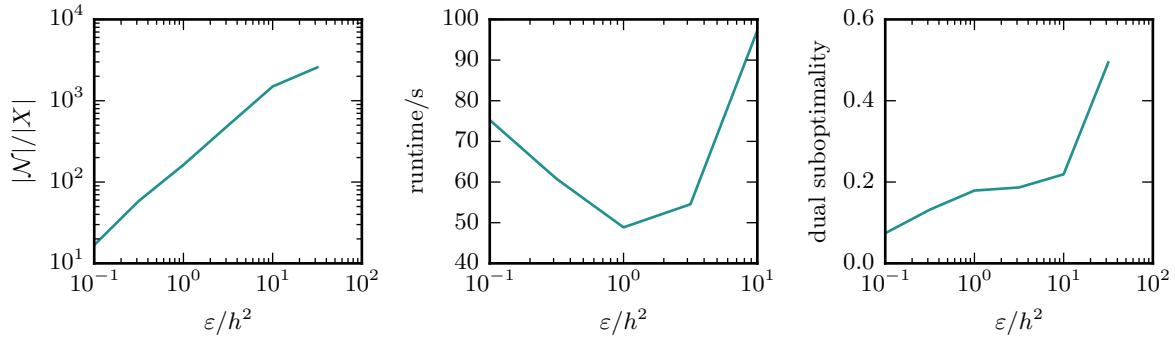

Figure 4: Different final values for ε in Algorithm 3.4. $X = Y$ are 2-d 256×256 grids. *Left* Number of variables in truncated kernel per $x \in X$. For $\varepsilon = 0.1 \cdot h^2$ only about 10 variables per $x \in X$ are required. As expected, this number increases with ε (cf. Fig. 1). *Center* For large ε, the runtime decreases with ε, since the number of variables decreases (cf. left plot). For smaller ε, the runtime increases again, since more stages of ε-scaling are required. *Right* The optimal regularized dual variables were transformed into feasible unregularized dual variables, by decreasing each $\alpha(x)$ until all dual constraints $\alpha(x) + \beta(y) \leq c(x, y)$ were met, (2.2b). The sub-optimality of these dual variables is shown. As expected (see Sect. 1.1) they converge towards a dual optimizer. Absolute optimal value was between 100 and 400 for the used test problems, i.e. for small ε, sub-optimality is small compared to total scale.

5.3 Versatility

The framework of scaling algorithms developed in [19], see Sect. 2, allows to solve more general transport-type problems for which the enhancements of Sect. 3 still apply. We now give some examples to demonstrate this flexibility. The scope of the following examples is similar to [19], but with Algorithm 3.4 one can solve larger problems with smaller regularization.

KL **Fidelity and Wasserstein-Fisher-Rao distance.** For the marginal function $F_X(\sigma) = \lambda \cdot \mathrm{KL}_X(\sigma|\mu)$ with a given reference measure $\mu \in \mathbb{R}_+^X$ and a weight $\lambda > 0$, see Def. 2.3, one obtains for the (stabilized) proxdiv operator

$$\mathrm{proxdiv}_\varepsilon F_X(\sigma) = (\mu \oslash \nu)^{\frac{\lambda}{\lambda+\varepsilon}}, \qquad \mathrm{proxdiv}_\varepsilon F_X(\sigma, \alpha) = \exp\left(-\tfrac{\alpha}{\lambda+\varepsilon}\right) \odot (\mu \oslash \nu)^{\frac{\lambda}{\lambda+\varepsilon}}. \tag{5.1}$$

A proof is given in [19]. Compared to the standard Sinkhorn algorithm, the only modification is the pointwise power of the iterates. As $\lambda \to \infty$ the Sinkhorn iterations are recovered. Note that in the stabilized operator only the exponential $\exp\left(\frac{-\alpha}{\lambda+\varepsilon}\right)$ needs to be evaluated, which remains bounded as $\varepsilon \to 0$. Algorithm 3.4 performs similarly with KL-fidelity as with fixed marginal constraints, allowing to efficiently solve large unbalanced transport problems. Since the truncation scheme can also be used with non-standard cost functions such as (2.7), this includes in particular the Wasserstein-Fisher-Rao (WFR) distance.

Fig. 5 shows two geodesics for the WFR distance, to intuitively illustrate its properties. The geodesics have been computed as weighted barycenters between their endpoints (see below). For a direct dynamic formulation we refer to [31, 17, 35]. For the relation to the KL soft-marginal formulation, Def. 2.3, see [35].

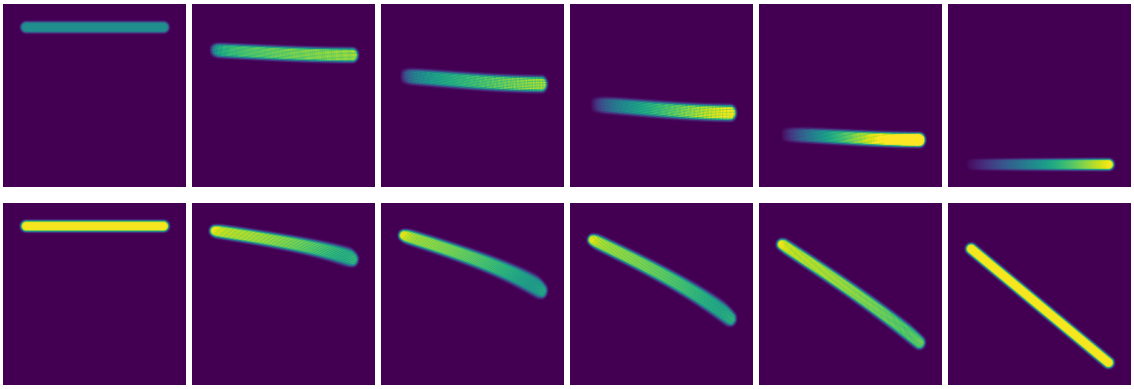

Figure 5: Geodesics for the Wasserstein-Fisher-Rao distance on $[0,1]^2$, approximated by a 256×256 grid, computed as barycenters between the end-points with varying weights. *Top row* Mass can increase or decrease during transport to match the target density. *Bottom row* Mass that must travel further, is decreased during transport to save cost.

Wasserstein barycenters. Wasserstein barycenters as a natural generalization of the Riemannian center of mass have been studied in [1]. The computation of entropy regularized Wasserstein barycenters with a Sinkhorn-type scaling algorithm has been presented in [7], an alternative numerical approach can be found in [24]. The iterations can be considered as a special case of the framework in [19]. Here, we very briefly recall the iterations. Derivations and proofs can be found in [7].

We want to compute the (entropy regularized) Wasserstein barycenter of a tuple $(\mu_1, \dots, \mu_n) \in \mathbb{R}^{X \times n}$ over a common base space $X = Y$ with metric d with non-negative weights $(\lambda_1, \dots, \lambda_n)$ that sum to one. The primal functional can be written as an optimization problem over a tuple $(\pi_i)_{i=1}^n = (\pi_1, \dots, \pi_n) \in \mathbb{R}^{(X \times X) \times n}$ of couplings, which requires a slight generalization of Def. 2.4, see [19]. It is given by

$$E((\pi_i)_i) = F_1((\mathrm{P}_X\, \pi_i)_i) + F_2((\mathrm{P}_Y\, \pi_i)_i) + \sum_{i=1}^n \lambda_i\, \mathrm{KL}(\pi_i | K) \tag{5.2a}$$

where

$$F_1((\nu_i)_i) = \sum_{i=1}^n \iota_{\{\mu_i\}}(\nu_i), \qquad F_2((\nu_i)_i) = \begin{cases} 0 & \text{if } \exists\, \sigma \in \mathbb{R}^X \text{ s.t. } [\sigma = \nu_i\, \forall\, i = 1, \dots, n], \\ +\infty & \text{else.} \end{cases}$$

and K is the kernel (2.8) over $X \times X$ for the cost $c = d^2$. When an optimizer $(\pi_i^\dagger)_i$ is found, the common second marginal of all $\pi_i^\dagger$ is the sought-after barycenter. A corresponding dual functional is

$$J((\alpha_i)_i, (\beta_i)_i) = -F_1^*(-(\lambda_i\, \alpha_i)_i) - F_2^*(-(\lambda_i\, \beta_i)_i) - \varepsilon \sum_{i=1}^n \lambda_i\, \mathrm{KL}^*([\mathrm{P}_X^\top\, \alpha_i + \mathrm{P}_Y^\top\, \beta_i]/\varepsilon | K) \tag{5.2b}$$

with

$$F_1^*((\alpha_i)_i) = \sum_{i=1}^n \langle \mu_i, \alpha_i \rangle\ , \qquad F_2^*((\beta_i)_i) = \begin{cases} 0 & \text{if } \sum_{i=1}^n \beta_i = 0, \\ +\infty & \text{else.} \end{cases} \tag{5.3}$$

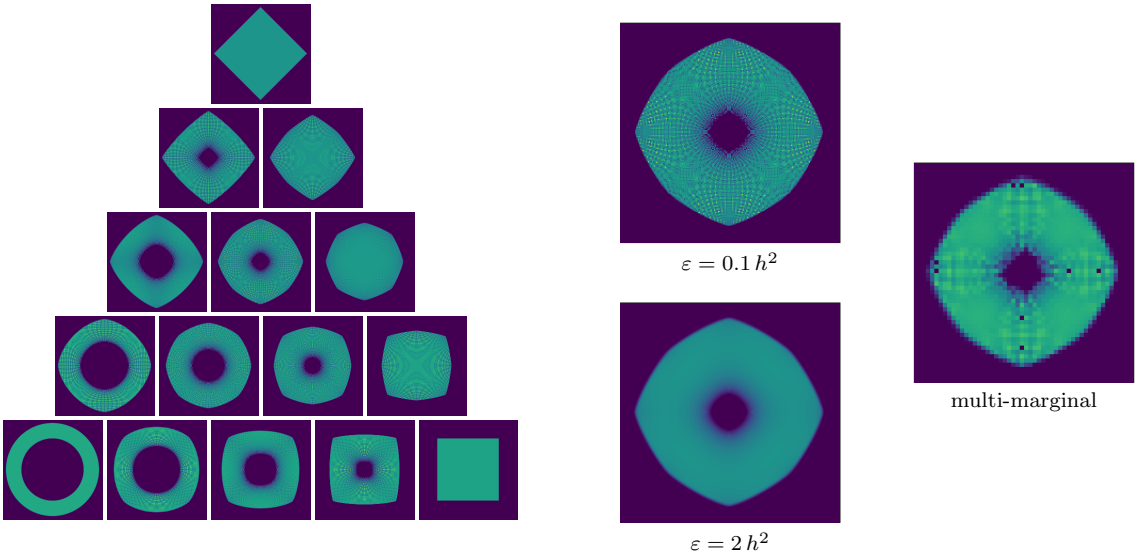

Figure 6: Barycenters in Wasserstein space over $[0,1]^2$, computed on 256×256 grids for $\varepsilon = 0.1\,h^2$. *Left* 'barycentric triangle' spanned by a ring, a diamond and a square (for weights see (5.5)). *Center* Close-up of the $\lambda = (1,2,1)/4$ barycenter, also shown for $\varepsilon = 2\,h^2$ (as reported in [7]). The $\varepsilon = 0.1\,h^2$ version is much sharper, revealing discretization artifacts. *Right* The same barycenter on a 64×64 grid, computed via multi-marginal formulation.

Note that relative to the dualization (2.10a) $\to$ (2.10b) we use rescaled dual variables $[(\alpha_i/\lambda_i, \beta_i/\lambda_i)]_i$. Primal and dual iterates are related by $\pi_i = \operatorname{diag}(\exp(\alpha_i/\varepsilon))K \operatorname{diag}(\exp(\beta_i/\varepsilon))$. As before, we introduce scaling factors $[(u_i, v_i)]_i$, $u_i = \exp(\alpha_i/\varepsilon)$, $v_i = \exp(\beta_i/\varepsilon)$ and optimize (5.2b) by alternating optimization in α and β, expressed by (u,v). Each α_i can be optimized independently and one gets a standard Sinkhorn update for every u_i: $u_i^{(\ell+1)} \overset{\text{def.}}{=} \mu_i \oslash (K\, v_i^{(\ell)})$. Optimization in $(\beta_i)_i$ is more involved. One finds $v_i^{(\ell+1)} \overset{\text{def.}}{=} \sigma^{(\ell+1)} \oslash (K^\top u_i^{(\ell+1)})$ where $\sigma^{(\ell+1)} = \prod_{i=1}^n (K^\top u_i^{(\ell+1)})^{\lambda_i}$ is the geometric mean of all $(K^\top u_i^{(\ell+1)})_i$, weighted by $(\lambda_i)_i$. Note that the prox and proxdiv steps of F_2 must be computed w.r.t. the weighted KL divergence

$$\mathrm{KL}^\lambda((\nu_i)_i|(\mu_i)_i) = \sum_{i=1}^n \lambda_i\, \mathrm{KL}(\nu_i|\mu_i)$$

to account for the weights $(\lambda_i)_i$ in the third term of the dual functional.

These iterations can be stabilized, analogous to Sect. 3.1, where one splits each (u_i, v_i) into pairs $(\tilde{u}_i, \tilde{v}_i)$ and $(\hat{\alpha}_i, \hat{\beta}_i)$ and introduces one stabilized kernel $\mathcal{K}_i$ per pair $(\hat{\alpha}_i, \hat{\beta}_i)$. The stabilized proxdiv step of F_2 w.r.t. KL^λ is given by

$$\operatorname{proxdiv}_{\varepsilon F_2}((\nu_i)_i, (\beta_i)_i) = \frac{\sigma}{\nu_i} \qquad \text{where} \qquad \sigma = \exp\left(\sum_{i=1}^n \lambda_i\,(\log \nu_i - \beta_i/\varepsilon)\right). \tag{5.4}$$

The terms β_i/ε in the expression for σ are seemingly unstable in the limit $\varepsilon \to 0$. However, the F_2^* term in the dual (5.2b) enforces $\sum_{i=0}^n \lambda_i\,\beta_i = 0$. Therefore, only the deviation from this constraint enters the exponential in (5.4). When one gradually approaches the optimal solution during ε-scaling, this deviation can be kept numerically small. Likewise, adaptive truncation applies to each kernel separately and a multi-scale coarse-to-fine approach can be used, as outlined in

Sects. 3.3-3.4. While this means that one must store multiple kernels, due to sparsity this is no issue. Often one already has to keep different kernels, when one uses different discrete spaces X_i for each marginal, to better capture the respective supports of the marginals μ_i.

In Fig. 6 a barycentric triangle is shown, computed with Algorithm 3.4, for weights $(\lambda_1, \lambda_2, \lambda_3)$ as follows:

$$\begin{array}{ccccccccc} & & & & (0,4,0)/4 & & & & \\ & & & (1,3,0)/4 & & (0,3,1)/4 & & & \\ & & (2,2,0)/4 & & (1,2,1)/4 & & (0,2,2)/4 & & \\ & (3,1,0)/4 & & (2,1,1)/4 & & (1,1,2)/4 & & (0,1,3)/4 & \\ (4,0,0)/4 & & (3,0,1)/4 & & (2,0,2)/4 & & (1,0,3)/4 & & (0,0,4)/4 \end{array} \tag{5.5}$$

The log-domain stabilization allows to reach a lower final regularization ε as for example in [7]. In fact, the regularization can be made so small, that discretization artifacts become detectable. We do not argue that this is visually more pleasing. However, it clearly gives a better approximation to the original unregularized barycenter problem and illustrates that with log-domain stabilization entropy regularized numerical methods can produce sharp results. We conjecture that there are more elaborate ways of removing the artifacts than the rather indiscriminate entropic blur. Truncation of the kernel is also efficient for the barycenter computation: On the finest hierarchy level, the average worst-case sparsity of the kernels was $3.5 \cdot 10^{-3}$, for the final ε it was $6.0 \cdot 10^{-4}$.

Wasserstein-Fisher-Rao barycenters. Similarly one can define barycenters for transport distances with KL marginal fidelity, which includes the Gaussian Hellinger-Kantorovich (GHK) distance and the Wasserstein-Fisher-Rao (WFR) distance (Def. 2.3). Primal and dual functional are given by (5.2) with

$$F_1((\nu_i)_i) = \Lambda \cdot \sum_{i=1}^{n} \lambda_i \, \mathrm{KL}(\nu_i|\mu_i)\,, \qquad F_2((\nu_i)_i) = \inf_{\sigma \in \mathbb{R}^X} \Lambda \cdot \sum_{i=1}^{n} \lambda_i \, \mathrm{KL}(\nu_i|\sigma)\,,$$

where $\Lambda > 0$ is a global weight of the KL-fidelity. When a primal optimizer is found, the minimizing σ in F_2 yields the sought-after barycenter. We refer to [19] for details. Partial optimization w.r.t α can again be done for each α_i separately, obtaining KL fidelity updates for each u_i, as given by (5.1). For the v-update one finds

$$\mathrm{proxdiv}_{\varepsilon} F_2((\nu_i)_i, (\beta_i)_i) = \nu_i^{\frac{-\Lambda}{\varepsilon+\Lambda}} \exp\left(-\tfrac{\beta_i}{\varepsilon+\Lambda}\right) \cdot \left(\sum_j \lambda_j \, \nu_j^{\frac{\varepsilon}{\varepsilon+\Lambda}} \exp\left(-\tfrac{\beta_j}{\varepsilon+\Lambda}\right) \right)^{\frac{\Lambda}{\varepsilon}}.$$

A barycentric triangle for the WFR distance is displayed in Fig. 7, a more detailed close-up with comparison to other models is shown in Fig. 8. Each reference measure consists of three spatially separated smaller objects with different mass. This implies that the Wasserstein barycenters have to transfer mass between the different groups, whereas the WFR barycenters can compensate the difference by creating or annihilating mass, resulting in more natural interpolations. Fig. 8 also shows a barycenter for the GHK distance, as computed with Gaussian convolution, without log-domain stabilization. Since the GHK uses the squared Euclidean distance as cost function, one can use efficient numerical methods for Gaussian convolution to avoid storing the dense kernel [49]. But this does not generalize to the cost function (2.7) for WFR

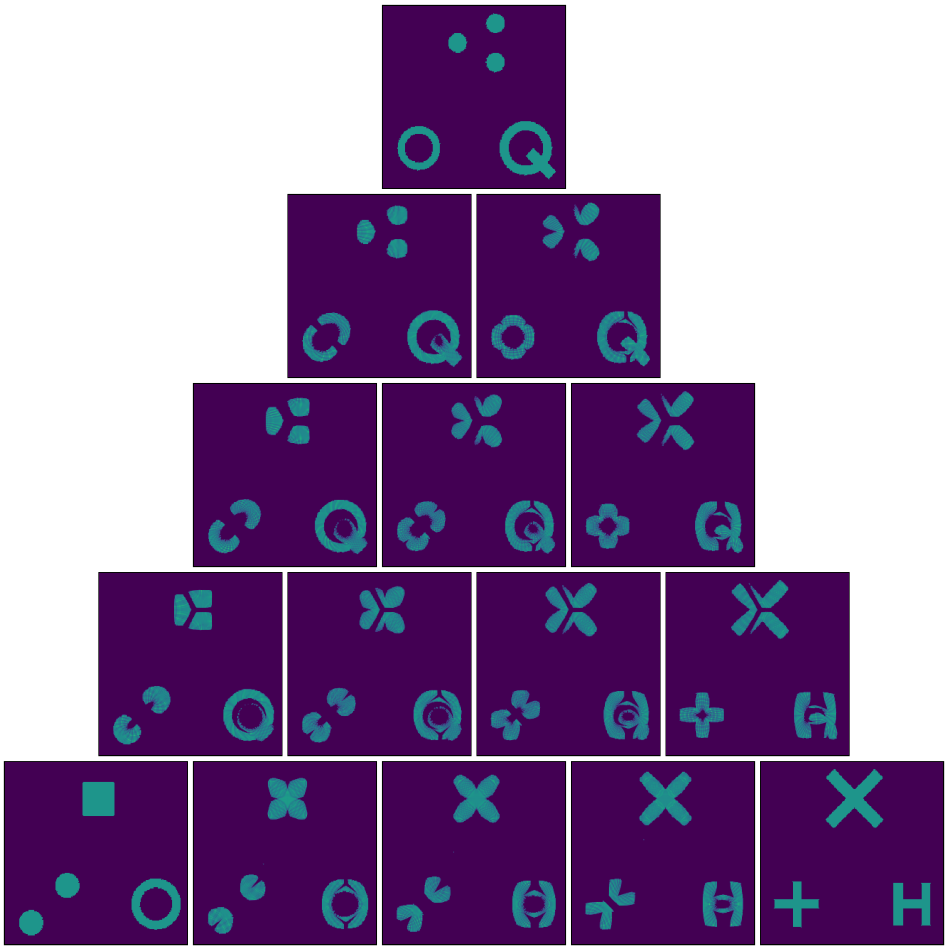

Figure 7: Entropic barycentric triangle for the Wasserstein-Fisher-Rao distance on $[0, \pi]^2$, approximated by a 256×256 grid, with $\varepsilon = 0.1\, h^2 \approx 8.5 \cdot 10^{-6}$, weights given by (5.5).

and neither can it be combined with log-domain stabilization. Since adaptive truncation of the kernel and multi-scale approximation of cost functions are rather flexible w.r.t. the cost function, Algorithm 3.4 can efficiently compute the WFR barycenter for small ε. Again, the barycenters computed with Algorithm 3.4 exhibit virtually no blur.

Multi-Marginal Problems. In [7] it was shown that the Sinkhorn algorithm can also be extended to multi-marginal problems and it is straightforward to adapt Algorithm 3.4 to this setting. For n marginals the kernel becomes an n-dimensional matrix and (α, β) are replaced by a list $(\alpha_1, \dots, \alpha_n)$. Stabilization is then done for each dual variable and for truncation in (3.10) one considers $\exp(-[c(x_1, \dots, x_n) - \sum_{i=1}^n \alpha_i(x_i)]/\varepsilon)$ where c is the n-dimensional cost function.

In [1] it is shown that barycenters in Wasserstein space can also be computed via a multi-marginal formulation. As a toy example we computed a weighted barycenter between three 64×64 images (where pixels with value zero were ignored). Projection of the optimal multi-dimensional coupling to the barycenter image is done as described in [7]. The result is shown in Fig. 6. Cardinality of the full product space was $1.5 \cdot 10^9$. The final sparsity of the truncated multi-dimensional kernel was $1.2 \cdot 10^{-3}$, worst-case sparsity was $4.7 \cdot 10^{-3}$.

Unfortunately, the dimensionality of multi-marginal problems grows so quickly with image

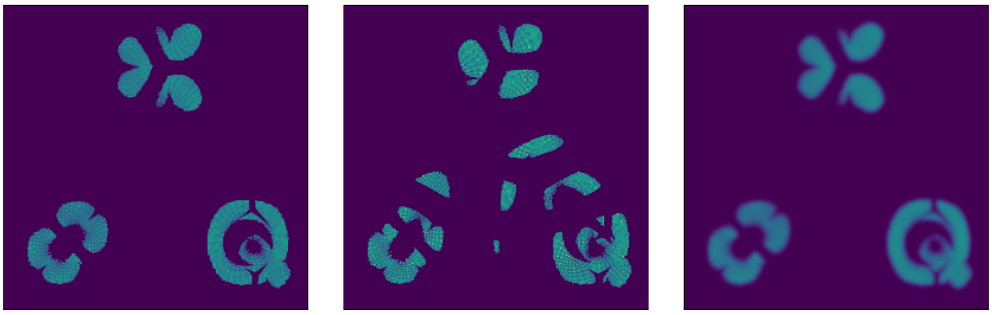

Figure 8: Comparison of different barycenters models: *Left* Close-up of Wasserstein-Fisher-Rao barycenter for $\lambda = (1,2,1)/4$ from Fig. 7, for $\varepsilon = 0.1\,h^2$. *Center* Corresponding Wasserstein barycenter between normalized reference measures for $\varepsilon = 0.1\,h^2$. *Right* Corresponding Gaussian Hellinger-Kantorovich barycenter for $\varepsilon = 6.55\,h^2$, as computed with Gaussian convolution without log-domain stabilization.

size, that at this point this does not yet scale efficiently to larger problems. For now, tricks to avoid storing the full kernel that are tailored to the structure of particular cost functions are more viable, see for example [7]. Nonetheless, we consider this an interesting proof of concept for the flexibility of Algorithm 3.4 and will further study this application.

Wasserstein Gradient Flows. In [39] diagonal scaling algorithms were extended to compute proximal steps for entropy regularized optimal transport to approximate gradient flows in Wasserstein space (cf. Sect. 1.1). This was then subsumed into the general framework of [19]. Here we given an example for the porous medium equation, for more details we refer to [39, 19]. Let

$$F : \mathcal{P}(X) \to \overline{\mathbb{R}}, \qquad \mu \mapsto \sum_{x\in X} u\left(\frac{\mu(x)}{\mathcal{L}(x)}\right) \mathcal{L}(x) + \sum_{x\in X} v(x)\,\mu(x) \tag{5.6}$$

where $u(s) = s^2$, $\mathcal{L}$ is the discretized Lebesgue measure on $X \subset \mathbb{R}^d$ and $v : X \to \overline{\mathbb{R}}$ is a potential. Then, for some initial $\mu^{(0)} \in \mathcal{P}(X)$ and a time step size $\tau > 0$ we iteratively construct a sequence $(\mu^{(\ell)})_\ell$ where $\mu^{(\ell+1)}$ is given by the proximal step of F with step size τ w.r.t. the entropy regularized Wasserstein distance on X from reference point $\mu^{(\ell)}$. Based on Def. 2.7, $\mu^{(\ell+1)}$ can be computed as follows:

$$\pi^{(\ell+1)} \overset{\text{def.}}{=} \operatorname*{argmin}_{\pi\in\mathcal{P}(X^2)} \left(\iota_{\{\mu^{(\ell)}\}}(\mathrm{P}_X\,\pi) + 2\,\tau\cdot F(\mathrm{P}_Y\,\pi) + \varepsilon\,\mathrm{KL}(\pi|K)\right), \quad \mu^{(\ell+1)} \overset{\text{def.}}{=} \mathrm{P}_Y \pi^{(\ell+1)}, \tag{5.7}$$

where K is the kernel w.r.t. the squared Euclidean distance on X. Then introduce the time-continuous interpolation

$$\overline{\mu} : \mathbb{R}_+ \to \mathcal{P}(X), \qquad t \mapsto \mu^{(\ell)} \quad \text{where} \quad t \in [\tau\cdot\ell, \tau\cdot(\ell+1)). \tag{5.8}$$

Consider now the limit $(\tau,\varepsilon) \to 0$ in a way such that $\varepsilon|\log\varepsilon| \le \tau^2$. Then, up to discretization, the function $\overline{\mu}$ converges to a solution of the porous media PDE

$$\partial_t \overline{\mu} = \Delta(\overline{\mu}^2) + \mathrm{div}(\overline{\mu}\cdot\nabla v)\,. \tag{5.9}$$

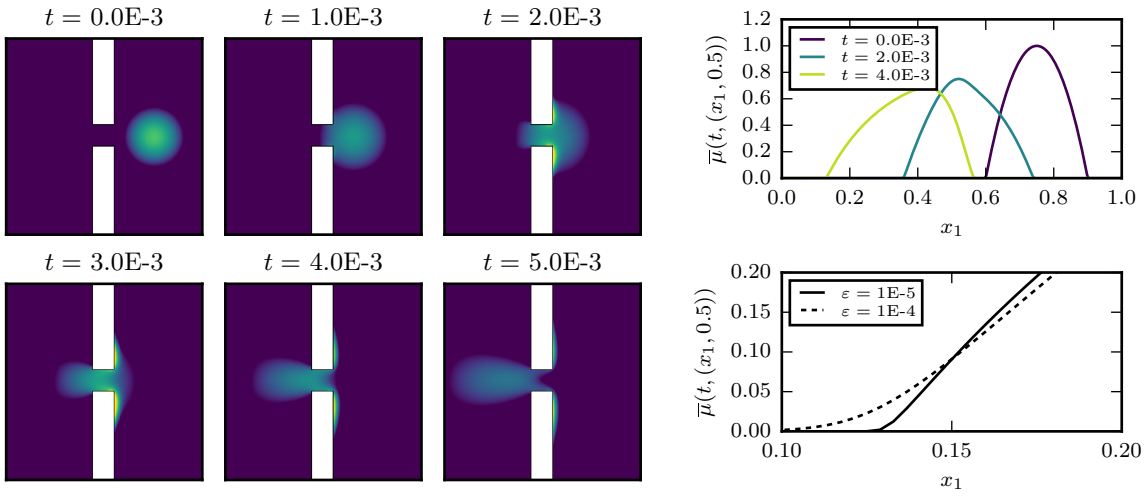

Figure 9: *Left* Entropic Wasserstein gradient flow for the porous media equation on $[0,1]^2$, approximated by a 256×256 grid with $\varepsilon = 10^{-5} \approx 0.66\,h^2$, $\tau = 2 \cdot 10^{-4}$. The energy is given by (5.6) with $v((x_1,x_2)) = 100 \cdot x_1$ if $x = (x_1,x_2) \in \Omega$, $v(x) = +\infty$ otherwise and $\Omega = [0,1]^2 \setminus \hat{\Omega}$ where $\hat{\Omega}$ is a 'barrier' indicated by the white rectangles. *Top Right* Cross section of density at different times along $x_2 = 0.5$. *Bottom Right* Close-up for $t = 4 \cdot 10^{-3}$ for different values of regularization ε. For $\varepsilon = 10^{-5}$ the compact support of $\overline{\mu}$, a characteristic feature of the porous media equation, is numerically well preserved. Without log-domain stabilization, for $\varepsilon = 10^{-4}$ the entropic blur quickly distorts this feature.

A proof is given in [16]. Problem (5.7) is an instance of Def. 2.4 and can be solved by alternating dual optimization. The updates w.r.t. $\iota_{\{\mu^{(\ell)}\}}$ are standard Sinkhorn iterations. The proximal step of F can be shown to be

$$\left[\mathrm{prox}_{\varepsilon F}(\nu)\right](x) = \left(\frac{\varepsilon}{4\,\tau}\mathcal{W}\left(\frac{4\,\tau}{\varepsilon}\cdot\frac{\mu(x)}{\mathcal{L}(x)}\cdot\exp(-v(x)/\varepsilon)\right)\right)\cdot\mathcal{L}(x)\,. \tag{5.10}$$

where $\mathcal{W}$ denotes the Lambert W function (or product logarithm). To avoid numerical issues, we propose to use the asymptotic expansion $\mathcal{W}(\exp(s)) = s - \log(s) + \log(s)/s + o(1)$ as $s \to +\infty$ [21] whenever $-v(x)/\varepsilon \gg 1$. The same applies for the stabilized iteration.

A numerical example is shown in Fig. 9. As in the previous experiments, Algorithm 3.4 allows to use log-domain stabilization on large problems, producing sharp results. In this example, the compact support of the porous media equation is numerically well preserved.

6 Conclusion

Scaling algorithms for entropy regularized transport-type problems have become a wide-spread numerical tool. Naive implementations have some severe numerical limitations, in particular for small regularization and on large problems. In this article, we proposed an enhanced variant of the standard scaling algorithm to address these issues: Diverging scaling factors and slow convergence are remedied by log-domain stabilization and ε-scaling. Required runtime and memory are significantly reduced by adaptive kernel truncation and a coarse-to-fine scheme. A new convergence analysis for the Sinkhorn algorithm was developed. Numerical examples showed the

efficiency of the enhanced algorithm, confirmed the scaling predicted by the convergence analysis and demonstrated that the algorithm can produce sharp results on a wide range of transport-type problems. Potential directions for future research are the more detailed study of ε-scaling, a more systematic understanding of the stability of the log-domain stabilization and application to larger multi-marginal problems.

Acknowledgements. Lénaïc Chizat, Luca Nenna and Gabriel Peyré are thanked for stimulating discussions. Bernhard Schmitzer was supported by the European Research Council (project SIGMA-Vision).

A Additional Proofs

A.1 Proof of Lemma 4.20

First, we establish existence of minimizers. For some $\varepsilon > 0$, $d \in \mathbb{R}^{R\times R}$ the functional $\beta \mapsto \hat{J}_{\varepsilon,d}(\beta)$ is convex and bounded from below. Further, it is invariant under adding the same constant to all components of β. Hence, the optimal value $\min_\beta \hat{J}_{\varepsilon,d}(\beta)$ is not changed by adding the constraint $\beta_1 = 0$. With this added constraint the functional becomes strictly convex and coercive in the remaining variables, hence a unique minimizer exists. The full set of minimizers is then obtained via constant shifts.

The first order optimality condition for the functional yields for the i-th component of β:

$$\beta(i) = \frac{1}{2}\left[\operatorname*{softmax}_{j:j\neq i}(-d(i,j)+\beta(j),\varepsilon) + \operatorname*{softmin}_{j:j\neq i}(d(j,i)+\beta(j),\varepsilon)\right],$$

where the subscript $j : j \neq i$ denotes that softmax is taken only over components $\{1,\dots,R\}\setminus\{i\}$. Finiteness of d ensures that this expression is meaningful. Let $i_1 \in \{1,\dots,R\}$ be an index where $\Delta\beta$ is maximal, i.e. $\Delta\beta(i_1) = \max \Delta\beta$.

From the optimality conditions for $\beta_a(i_1)$, $a = 1, 2$, and (1.5) we obtain:

$$\begin{aligned}
\beta_a^\dagger(i_1) &= \frac{1}{2}\left[\operatorname*{softmax}_{j:j\neq i_1}(-d_a(i_1,j)+\beta_a^\dagger(j),\varepsilon_a) + \operatorname*{softmin}_{j:j\neq i_1}(d_a(j,i_1)+\beta_a^\dagger(j),\varepsilon_a)\right],\\
\Delta\beta(i_1) &\leq \frac{1}{2}\left[\max_{j:j\neq i_1}(-\Delta d(i_1,j)+\Delta\beta(j)) + \max_{j:j\neq i_1}(\Delta d(j,i_1)+\Delta\beta(j)) + (\varepsilon_1+\varepsilon_2)\cdot\log R\right]\\
&\leq \max_{j:j\neq i_1}(w(i_1,j)+\Delta\beta(j)) + \tfrac{\varepsilon_1+\varepsilon_2}{2}\log R
\end{aligned}$$

where $w(i,j) = \max\{-\Delta d(i,j), \Delta d(j,i)\}$. This implies there is some $i_2 \in \{1,\dots,R\}\setminus\{i_1\}$ with

$$\Delta\beta(i_2) \geq -w(i_1,i_2) + \Delta\beta(i_1) - \tfrac{\varepsilon_1+\varepsilon_2}{2}\cdot\log R\,.$$

We will call the index i_2 a child of i_1. We now repeat this reasoning to derive lower bounds for other entries of $\Delta\beta$. For this we must 'remove' the index i_2 from the problem, defining a reduced problem. Let $I_1 = \{i_1, i_2\}$ and let $I_2 = \{1,\dots,R\}\setminus I_1$. We will keep all variables of β with indices in I_2, but describe all variables with indices in I_1 by a single reduced variable. For this we consider vectors in $\mathbb{R}^{1+|I_2|}$, where we index the entries by $\{i_1\}\cup I_2$. One can think of this as a vector in $\mathbb{R}^R$, where we have 'crossed out' entries corresponding to I_1 and replaced them by a single effective entry, indexed with i_1. For $a = 1, 2$ we consider the reduced functionals

$$\hat{\mathcal{J}}_a : \hat{\beta} \mapsto \hat{J}_{\varepsilon_a,d_a}(\tilde{\beta}_a + B\,\hat{\beta})$$

where $\tilde\beta_a \in \mathbb{R}^R$ is a constant offset, $\hat\beta \in \mathbb{R}^{1+|I_2|}$ is the reduced variable and $B \in \mathbb{R}^{R\times(1+|I_2|)}$ is a matrix that implements the parametrization. We set

$$\tilde\beta_a(j) = \begin{cases} \beta_a^\dagger(j) - \beta_a^\dagger(i_1) & \text{if } j \in I_1, \\ 0 & \text{else,} \end{cases} \qquad B(j,k) = \begin{cases} 1 & \text{if } j \in I_1,\ k = i_1, \\ 1 & \text{if } j = k \in I_2, \\ 0 & \text{else.} \end{cases}$$

So the reduced functionals are given by

$$\begin{aligned} \hat{\mathcal{J}}_a(\hat\beta) = &\sum_{\substack{j\in I_1,\\ k\in I_1}} \exp([-d_a(j,k) - \tilde\beta_a(j) + \tilde\beta_a(k)]/\varepsilon_a) \\ &+ \sum_{\substack{j\in I_1,\\ k\in I_2}} \exp([-d_a(j,k) - \tilde\beta_a(j) - \hat\beta(i_1) + \hat\beta(k)]/\varepsilon_a) \\ &+ \sum_{\substack{j\in I_2,\\ k\in I_1}} \exp([-d_a(j,k) - \hat\beta(j) + \tilde\beta_a(k) + \hat\beta(i_1)]/\varepsilon_a) \\ &+ \sum_{\substack{j\in I_2,\\ k\in I_2}} \exp([-d_a(j,k) - \hat\beta(j) + \hat\beta(k)]/\varepsilon_a) \\ = &\sum_{\substack{j\in \{i_1\}\cup I_2,\\ k\in \{i_1\}\cup I_2}} \exp([-D_a(j,k) - \hat\beta(j) + \hat\beta(k)]/\varepsilon_a) \end{aligned}$$

with the reduced coefficient matrix $D_a \in \mathbb{R}^{(1+|I_2|)^2}$ with entries

$$D_a(j,k) = \begin{cases} \operatorname{softmin}_{r\in I_1,\, s\in I_1} \left(d_a(r,s) + \tilde\beta_a(r) - \tilde\beta_a(s), \varepsilon_a\right) & \text{if } j = i_1,\ k = i_1, \\ \operatorname{softmin}_{r\in I_1} \left(d_a(r,k) + \tilde\beta_a(r), \varepsilon_a\right) & \text{if } j = i_1,\ k \in I_2, \\ \operatorname{softmin}_{s\in I_1} \left(d_a(j,s) - \tilde\beta_a(s), \varepsilon_a\right) & \text{if } j \in I_2,\ k = i_1, \\ d_a(j,k) & \text{if } j \in I_2,\ k \in I_2. \end{cases}$$

Consider the reduced variables $\hat\beta_a^\dagger \in \mathbb{R}^{1+|I_2|}$ with entries

$$\hat\beta_a^\dagger(j) = \begin{cases} \beta_a^\dagger(i_1) & \text{if } j = i_1, \\ \beta_a^\dagger(j) & \text{if } j \in I_2. \end{cases}$$

Then $\beta_a^\dagger = \tilde\beta_a + B\,\hat\beta_a^\dagger$ and therefore $\hat\beta_a^\dagger$ are minimizers of $\hat{\mathcal{J}}_a$. Note also that $\hat\beta_2^\dagger(j) - \hat\beta_1^\dagger(j) = \Delta\beta^\dagger(j)$ for $j \in \{i_1\} \cup I_2$.

Using the optimality conditions for the reduced functionals and arguing as above, we find

$$\Delta\beta(i_1) \le \max_{k\in I_2}(W(i_1,k) + \Delta\beta(k)) + \tfrac{\varepsilon_1+\varepsilon_2}{2} \log R$$

where $W(i_1,k) = \max\{-\Delta D(i_1,k), \Delta D(k,i_1)\}$ for $k \in I_2$ and $\Delta D = D_2 - D_1$. With (1.5) we find

$$\begin{aligned} -\Delta D(i_1,k) &\le \max_{j\in I_1}(-\Delta d(j,k) - \Delta\beta(j) + \Delta\beta(i_1)) + \varepsilon_2 \log R\,, \\ \Delta D(k,i_1) &\le \max_{j\in I_1}(\Delta d(k,j) - \Delta\beta(j) + \Delta\beta(i_1)) + \varepsilon_1 \log R \end{aligned}$$

and eventually $W(i_1,k) \le \max_{j\in I_1}(w(j,k)-\Delta\beta(j)) + \Delta\beta(i_1) + \max\{\varepsilon_1,\varepsilon_2\}\cdot\log R$. So there is some index $i_3 \in I_2$ such that

$$\Delta\beta(i_3) \ge \min_{j\in I_1}(-w(j,i_3)+\Delta\beta(j)) - \left[\tfrac{\varepsilon_1+\varepsilon_2}{2} + \max\{\varepsilon_1,\varepsilon_2\}\right]\log R\,.$$

The index i_3 will be called a child of the minimizing index $j \in I_1$ on the r.h.s., (or one of the minimizing indices). Then we add i_3 to the set I_1 and repeat the argument with the reduced functional, to obtain an index i_4 and repeat this until I_1 contains all indices.

Since we assign every new index i_k that is added to I_1 as a child to one parent node in I_1, this also constructs a tree graph with root node i_1 (finiteness of d and consequently D implies that this graph is connected). For an index i_k let $(i_1, i_2, \dots, i_k)$ be the unique path from the root to i_k. Then

$$\begin{aligned}\Delta\beta(i_k) &\ge -\sum_{j=2}^{k} w(i_{j-1},i_j) + \Delta\beta(i_1) - 2\,(k-1)\,\varepsilon_1\,\log R\\ &\ge -\operatorname{maxdiam}(w) + \Delta\beta(i_1) - 2\,\varepsilon_1\,R\,\log R\,.\end{aligned}$$

Since $\Delta\beta(i_1) = \max\Delta\beta$ the result follows.

A.2 Proof of Theorem 4.16

Let π_1, π_2 be the primal optimizers associated with (α_1,β_1) and (α_2,β_2) and consider the assignment graph for π_1 and π_2 and threshold $1/M$ (see Lemma 4.18). Let $\{(X_i,Y_i)\}_{i=1}^R$ be the strongly connected components of the assignment graph. By virtue of Lemma 4.18, statement iii, $\mu(X_i) = \nu(Y_i)$ for $i = 1,\dots,R$. Pick some representatives $\{y_i\}_{i=1}^R \subset Y$ such that $y_i \in Y_i$ for $i = 1,\dots,R$.

For $a = 1, 2$, let now $\hat{J}_a$ be the reduced effective diagonal functionals, defined in Lemma 4.19, corresponding to spaces (X,Y), marginals (μ,ν), parameters ε_a, cost c, the partitions given by the strongly connected components and the representatives $\{y_i\}_{i=1}^R$. Let d_a be the corresponding effective coefficients (finite, since c is finite), let $\hat{\beta}_a^\dagger$ be two corresponding maximizers and let $\Delta d = d_2 - d_1$, $\Delta\hat{\beta} = \hat{\beta}_2^\dagger - \hat{\beta}_1^\dagger$. By virtue of Lemma 4.20 one has $\max\Delta\hat{\beta} - \min\Delta\hat{\beta} \le \operatorname{maxdiam}(w) + 2\,\varepsilon^1\,R\log R$, where $w \in \mathbb{R}^{R\times R}$ with $w(i,j) = \max\{-\Delta d(i,j), \Delta d(j,j)\}$.

Now we derive some estimates on Δd. Consider once more the assignment graph for π_1, π_2 and threshold $1/M$. For every edge $y \to x$ we have (using (2.11))

$$\alpha_1(x) + \beta_1(y) - c(x,y) \ge -\varepsilon_1\,\log M\,.$$

Moreover, from the marginal conditions we find $\pi_2(x,y) \le \nu(y)$, which implies

$$\alpha_2(x) + \beta_2(y) - c(x,y) \le \varepsilon_2\,\log M.$$

Combining the two estimates, we obtain

$$\Delta\alpha(x) + \Delta\beta(y) \le (\varepsilon_1+\varepsilon_2)\,\log M \overset{\text{def.}}{=} L$$

Similarly, for edges $x \to y$ we obtain

$$\Delta\alpha(x) + \Delta\beta(y) \ge -(\varepsilon^1+\varepsilon^2)\,\log M = -L\,.$$

Let now $(y_1, x_1, \dots, y_k)$ be an alternating path in (X, Y), then, by combining the above inequalities we find $\Delta\beta(y_{j+1}) \geq \Delta\beta(y_j) - 2 \cdot L$ for $j = 1, \dots, k-1$ and eventually

$$\Delta\beta(y_k) - \Delta\beta(y_1) \geq -2 \cdot (k-1) \cdot L\,.$$

Similarly, for a path $(\hat{x}_1, \hat{y}_2, \dots, \hat{y}_k)$ get

$$\Delta\alpha(x_1) + \Delta\beta(y_k) \geq -(2\,k-1) \cdot L\,,$$

and for a path $(\hat{y}_1, \hat{x}_1, \dots, \hat{y}_k, \hat{x}_k)$

$$\Delta\alpha(x_k) + \Delta\beta(y_1) \leq (2\,k-1) \cdot L.$$

Consider now a partition cell (X_i, Y_i) and let $y_i \in Y_i$ be the selected 'representative', as described above. For every $y \in Y_i$ there is a path to and from y_i with at most $2(|Y_i| - 1)$ edges, for every $x \in X_i$ there is a path to and from y_i with at most $2\,|Y_i| - 1$ edges. With $\Delta\tilde{\alpha}(x) = \Delta\alpha(x) + \Delta\beta(y_i)$ and $\Delta\tilde{\beta}(y) = \Delta\beta(y) - \Delta\beta(y_i)$ we therefore obtain

$$|\Delta\tilde{\alpha}(x)| \leq (2\,|Y_i| - 1) \cdot L, \qquad\qquad |\Delta\tilde{\beta}(y)| \leq 2\,(|Y_i| - 1) \cdot L.$$

We recall (4.7)

$$d_a(i,j) = \operatorname*{softmin}_{\substack{x \in X_i,\\ y \in Y_j}} \Big(c(x,y) - \tilde{\alpha}_a(x) - \tilde{\beta}_a(y) - \varepsilon_a \, \log(\mu(x)\,\nu(y)), \varepsilon_a \Big)$$

and get

$$\begin{aligned}
\Delta d(i,j) &\leq \max_{\substack{x \in X_i,\\ y \in Y_j}} \Big(-\Delta\tilde{\alpha}(x) - \Delta\tilde{\beta}(y) - \Delta\varepsilon \cdot \log\big(\mu(x)\,\nu(y)\big) \Big) + \varepsilon_1 \cdot \log\big(|X_i|\,|Y_j|\big) \\
&\leq 4\,|Y_i|\,L + \varepsilon_1 \cdot \log\big(|X_i|\,|Y_j|\big) \\
\Delta d(i,j) &\geq \min_{\substack{x \in X_i,\\ y \in Y_j}} \Big(-\Delta\tilde{\alpha}(x) - \Delta\tilde{\beta}(y) - \Delta\varepsilon \cdot \log\big(\mu(x)\,\nu(y)\big) \Big) - \varepsilon_2 \cdot \log\big(|X_i|\,|Y_j|\big) \\
&\geq -4\,|Y_i|\,L - \varepsilon_2 \cdot \log\big(|X_i|\,|Y_j|\big)
\end{aligned}$$

From this follows that $w(i,j) \leq 8 \max\{|Y_i|, |Y_j|\}\,\varepsilon_1 \log M + 2\,\varepsilon_1 \log N$, which in turn implies that $\operatorname{maxdiam}(w) \leq 16\,\varepsilon_1\,N \log M + 2\,\varepsilon_1\,R \log N$.

Recall that $\Delta\hat{\beta} = \hat{\beta}_2^\dagger - \hat{\beta}_1^\dagger$, where $\hat{\beta}_a^\dagger$, $a = 1, 2$, are the optimizers of the effective diagonal problems. Then from Lemma 4.20, and by bounding $R \leq N$ we obtain that

$$\max \Delta\hat{\beta} - \min \Delta\hat{\beta} \leq \varepsilon_1\,N\,(4 \log N + 16 \log M)$$

and finally with $\max \Delta\beta - \min \Delta\beta \leq \max \Delta\tilde{\beta} - \min \Delta\tilde{\beta} + \max \Delta\hat{\beta} - \min \Delta\hat{\beta}$ we get

$$\max \Delta\beta - \min \Delta\beta \leq \varepsilon_1\,N\,(4 \log N + 24 \log M),$$

and analogously we get the equivalent bound for $\Delta\alpha$.

References

[1] M. Agueh and G. Carlier. Barycenters in the Wasserstein space. 43(2):904–924, 2011.

[2] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Inc., 1993.

[3] L. Ambrosio and N. Gigli. A user's guide to optimal transport, 2011.

[4] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics. Birkhäuser Boston, 2005.

[5] H. H. Bauschke and P. L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. CMS Books in Mathematics. Springer, 1st edition, 2011.

[6] J.-D. Benamou and Y. Brenier. A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.

[7] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré. Iterative Bregman projections for regularized transportation problems. 37(2):A1111–A1138, 2015.

[8] J.-D. Benamou, F. Collino, and J.-M. Mirebeau. Monotone and consistent discretization of the monge-ampere operator. https://arxiv.org/abs/1409.6694.

[9] J.-D. Benamou, B. D. Froese, and A. M. Oberman. Numerical solution of the optimal transportation problem using the Monge–Ampère equation. *Journal of Computational Physics*, 260(1):107–126, 2014.

[10] D. P. Bertsekas. A distributed algorithm for the assignment problem. Technical report, Lab. for Information and Decision Systems Report, MIT, May 1979.

[11] D. P. Bertsekas. The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of Operations Research*, 14:105–123, 1988.

[12] D. P. Bertsekas and J. Eckstein. Dual coordinate step methods for linear network flow problems. *Mathematical Programming, Series B*, 42:203–243, 1988.

[13] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[14] Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. 44(4):375–417, 1991.

[15] M. Burger, M. Franek, and C.-B. Schönlieb. Regularised regression and density estimation based on optimal transport. *Applied Mathematics Research eXpress*, 3 2012.

[16] G. Carlier, V. Duval, G. Peyré, and B. Schmitzer. Convergence of entropic schemes for optimal transport and gradient flows. http://arxiv.org/abs/1512.02783, 2015.

[17] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard. An interpolating distance between optimal transport and Fisher-Rao. http://arxiv.org/abs/1506.06430, 2015.

[18] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard. Unbalanced optimal transport: Geometry and Kantorovich formulation. http://arxiv.org/abs/1508.05216, 2015.

[19] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard. Scaling algorithms for unbalanced transport problems. http://arxiv.org/abs/1607.05816, 2016.

[20] R. Cominetti and J. San Martin. Asymptotic analysis of the exponential penalty trajectory in linear programming. *Mathematical Programming*, 67:169–187, 1992.

[21] R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the LambertW function. *Advances in Computational Mathematics*, 5(1):329–359, 1996.

[22] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transportation distances. pages 2292–2300, 2013. http://arxiv.org/abs/1306.0895.

[23] M. Cuturi and D. Avis. Ground metric learning. 15:533–564, 2014.

[24] M. Cuturi and A. Doucet. Fast computation of wasserstein barycenters. In *International Conference on Machine Learning*, 2014.

[25] J. H. Fitschen, F. Laus, and G. Steidl. Transport between RGB images motivated by dynamic optimal transport. *Journal of Mathematical Imaging and Vision*, 2016.

[26] J. Franklin and J. Lorenz. On the scaling of multidimensional matrices. *Linear Algebra and its Applications*, 114–115:717–735, 1989.

[27] A. V. Goldberg and R. E. Tarjan. Finding minimum-cost circulations by successive approximation. *Math. Oper. Res.*, 15(3):430–466, 1990.

[28] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent. Optimal mass transport for registration and warping. 60(3):225–240, December 2004.

[29] R. Jordan, D. Kinderlehrer, and F. Otto. The variational formulation of the Fokker-Planck equation. 29(1):1–17, 1998.

[30] P. A. Knight. The Sinkhorn-Knopp algorithm: Convergence and applications. *SIAM. J. Matrix Anal. & Appl.*, 30(1):261–275, 2008.

[31] S. Kondratyev, L. Monsaingeon, and D. Vorotnikov. A new optimal transport distance on the space of finite Radon measures. http://arxiv.org/abs/1505.07746, 2015.

[32] J. Kosowsky and A. Yuille. The invisible hand algorithm: Solving the assignment problem with statistical physics. *Neural Networks*, 7(3):477–490, 1994.

[33] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics*, 2:83–97, 1955.

[34] C. Léonard. From the Schrödinger problem to the Monge–Kantorovich problem. *Journal of Functional Analysis*, 262(4):1879–1920, 2012.

[35] M. Liero, A. Mielke, and G. Savaré. Optimal entropy-transport problems and a new Hellinger-Kantorovich distance between positive measures. http://arxiv.org/abs/1508.07941, 2015.

[36] J. Maas, M. Rumpf, C. Schönlieb, and S. Simon. A generalized model for optimal transport of images including dissipation and density modulation. submitted, 2014.

[37] Q. Mérigot. A multiscale approach to optimal transport. *Computer Graphics Forum*, 30(5):1583–1592, 2011.

[38] A. M. Oberman and Y. Ruan. An efficient linear programming method for optimal transportation. http://arxiv.org/abs/1509.03668, 2015.

[39] G. Peyré. Entropic approximation of Wasserstein gradient flows. 8(4):2323–2351, 2015.

[40] J. Rabin, G. Peyré, and L. D. Cohen. Geodesic shape retrieval via optimal mass transport. pages 771–784, 2010.

[41] J. Rabin and N. Papadakis. Convex color image segmentation with optimal transport distances. In *Scale Space and Variational Methods (SSVM 2015)*, pages 256–268, 2015.

[42] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover's distance as a metric for image retrieval. 40(2):99–121, 2000.

[43] F. Santambrogio. *Optimal Transport for Applied Mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and Their Applications*. Birkhäuser Boston, 2015.

[44] B. Schmitzer. A sparse multi-scale algorithm for dense optimal transport. *Journal of Mathematical Imaging and Vision*, 56(2):238–259, 2016.

[45] B. Schmitzer and C. Schnörr. A hierarchical approach to optimal transport. In *Scale Space and Variational Methods (SSVM 2013)*, pages 452–464, 2013.

[46] B. Schmitzer and C. Schnörr. Globally optimal joint image segmentation and shape matching based on Wasserstein modes. *Journal of Mathematical Imaging and Vision*, 52(3):436–458, 2015.

[47] M. Sharify, S. Gaubert, and L. Grigori. Solution of the optimal assignment problem by diagonal scaling algorithms. http://arxiv.org/abs/1104.3830v2, 2013.

[48] R. D. Sinkhorn and P. J. Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific J. Math*, 21(2):343—-348, 1967.

[49] J. Solomon, F. de Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas. Convolutional Wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Transactions on Graphics (Proc. SIGGRAPH 2015)*, 34(4):66:1–66:11, 2015.

[50] J. Solomon, A. Nguyen, A. Butscher, M. Ben-Chen, and L. Guibas. Soft maps between surfaces. *Computer Graphics Forum*, 31(5), 2012.

[51] J. Solomon, R. Rustamov, L. Guibas, and A. Butscher. Wasserstein propagation for semi-supervised learning. In *International Conference on Machine Learning*, 2014.

[52] M. Thorpe, S. Park, S. Kolouri, G. K. Rohde, and D. Slepčev. A transportation Lp distance for signal analysis. https://arxiv.org/abs/1609.08669, 2016.

[53] C. Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2009.

[54] W. Wang, D. Slepčev, S. Basu, J. A. Ozolek, and G. K. Rohde. A linear optimal transportation framework for quantifying and visualizing variations in sets of images. 101:254–269, 2012.